Mon. Not. R. Astron. Soc. **000**, 1 (2009) Printed 12 January 2013 (MN LaTeX style file v2.2)

arXiv:1102.1752v2 [astro-ph.CO] 15 Aug 2011

On Estimating Lyα Forest Correlations between Multiple Sightlines

Matthew McQuinn*[1] and Martin White[1,2]

[1] *Department of Astronomy, University of California, Berkeley, CA 94720*
[2] *Department of Physics, University of California, Berkeley, CA 94720*

12 January 2013

ABSTRACT

The next frontier of Lyα forest studies is the reconstruction of 3D correlations from a dense sample of background sources. The measurement of 3D correlations has the potential to improve constraints on fundamental cosmological parameters, ionizing background models, and the reionization history. This study addresses the sensitivity of spectroscopic surveys to 3D correlations in the Lyα forest. We show that the sensitivity of a survey to this signal can be quantified by just a single number, a noise-weighted number density of sources on the sky. We investigate how the sensitivity of a spectroscopic quasar (or galaxy) survey scales as a function of its depth, area, and redshift. We propose a simple method for weighting sightlines with varying S/N levels to estimate the correlation function, and we show that this estimator generally performs nearly as well as the minimum variance quadratic estimator. In addition, we show that the sensitivity of a quasar survey to the flux correlation function is generally maximized if it observes each field just long enough to achieve $S/N \approx 2$ in a 1 Å pixel on an L_* quasar while acquiring spectra for all quasars with $L > L_*$: Little is gained by integrating longer on the same targets or by including fainter quasars. We quantify how these considerations relate to constraints on the angular diameter distance, the curvature of space-time, and the reionization history.

Key words: cosmology: theory – cosmology: large-scale structure – quasars: absorption lines – intergalactic medium

1 INTRODUCTION

The Lyα forest is an established tool for studying structure formation, the intergalactic medium (IGM), and cosmological parameters. It has been used to place constraints on the linear-theory density correlations at ~ 1 comoving Mpc separations, smaller separations than other large-scale structure probes (Weinberg et al. 1997; Croft et al. 1998; McDonald et al. 2000; Seljak et al. 2005; Viel et al. 2005). In addition, it is our best tool for studying the thermal history of the intergalactic gas (Schaye et al. 2000; McDonald et al. 2001; Lidz et al. 2010; Becker et al. 2011), and it has been used to place a lower bound on the redshift of reionization (Becker et al. 2001; Fan et al. 2006). However, past Lyα forest analyses have only utilized correlations in the transmission within a single sightline (e.g., McDonald et al. 2005) or between a small number of sightlines (e.g., Williger et al. 2000; Liske et al. 2000; Hennawi & Prochaska 2007). This approach was justified because of the low sky density of quasars in previous wide-field surveys. However, the next generation of surveys will achieve densities of $\approx 10-100$ quasars per deg^2, allowing $S/N > 1$ on tens of comoving Mpc 3D modes. The ongoing Baryon Oscillation Spectroscopic Survey (BOSS) on the 2.5 m Sloan Telescope aims to measure 3D correlations from 1.6×10^5 quasars over 8000 deg^2. A similar survey but on a 4 m telescope, BigBOSS, is anticipated to begin in 2016.[1] In general, piggybacking a 3D Lyα survey on an optical spectroscopic survey comes with little cost owing to the low sky density of $z > 2$ quasars.

The advantages of 3D Lyα forest correlations have been enumerated in several previous theoretical studies. The traditional approach utilizing purely line-of-sight measurements will always be more sensitive to parameters that impact correlations on less than several Mpc scales. However,

---

[1] `http://cosmology.lbl.gov/BOSS/`, `http://bigboss.lbl.gov/index.html`

* mmcquinn@berkeley.edu

© 2009 RAS

unlike with line-of-sight measurements, with 3D correlation measurements it is possible to detect the baryon acoustic oscillation feature at high precision and to use it as a standard ruler (White 2003; McDonald & Eisenstein 2007; Slosar et al. 2009; White et al. 2010). Such a measurement at $z = 2 - 3$ with a BOSS-like spectroscopic quasar survey has the potential to rule out early dark energy models and to place 10^{-3}–level constraints on Ω_k, the $z = 0$ spatial curvature density (McDonald & Eisenstein 2007). Three-dimensional measurements are also more sensitive to large-scale intensity and temperature fluctuations than are line-of-sight ones (White et al. 2010; McQuinn et al. 2010). In fact, the resulting temperature fluctuations from physically motivated models for He II reionization can change the Lyα forest correlation function by $\mathcal{O}(1)$ at 100 Mpc separations. This difference would be easily detected in a measurement utilizing correlations between sightlines, but would have gone undetected in previous line-of-sight analyses (McQuinn et al. 2010). In addition, a future Lyα forest survey that is sensitive to $z \approx 4$ could detect the relic temperature fluctuations from hydrogen reionization (McQuinn et al. 2010).

This paper studies the ability of spectroscopic quasar surveys to constrain correlations in the 3D Lyα forest. Section 2 presents formulae for the sensitivity of spectroscopic surveys to 3D correlations as well as for how to weight sightlines based on their signal-to-noise properties. We use the resulting expressions to understand the sensitivity of both purely hypothetical and proposed quasar surveys as a function of depth, volume, and redshift. Section 3 discusses how the suggested weights relate to the minimum variance quadratic estimator. Section 4 discusses additional survey considerations: (1) whether taking the spectra of high-redshift galaxies in addition to quasars can enhance the sensitivity, (2) the advantages of cross correlating a 3D Lyα survey with another survey of large-scale structure, (3) how to optimize a survey's strategy to minimize the uncertainty in its estimate of the correlation function, and (4) how the previous considerations translate to constraints on cosmological parameters and interesting astrophysical processes. Finally, Section 5 discusses the severity of continuum subtraction errors, of errors in the mean flux estimate, and of contamination from damping wings. We find that these systematics are likely to be less severe for upcoming 3D Lyα forest analyses than they were for past line-of-sight studies.

We assume a flat ΛCDM cosmological model where necessary with $\Omega_m = 0.27$, $h = 0.71$, $\sigma_8 = 0.8$, $n_s = 0.96$, and $\Omega_b = 0.046$ (Komatsu et al. 2011). We henceforth will use "Mpc" as shorthand for "comoving Mpc," and we use the standard Fourier convention for cosmological studies in which $2\,\pi$'s appear under the dk's. Table 1 provides relevant numbers for the Lyα forest that are used in our calculations.

Table 1. Relevant numbers and conversion factors: χ is the conformal distance, $\Delta\chi$ is the conformal distance covered by the Lyα forest that is not contaminated by the Lyβ forest, and b is the large-scale bias of the Lyα forest in our model. The mean flux, $\langle F \rangle$, is calculated using the fitting formula in Meiksin (2009), and $d\lambda/d\chi$ and $dv/d\chi$ are in units of Å Mpc^{-1} and km s^{-1} Mpc^{-1}, respectively.

z	λ [Å]	$\langle F \rangle$	$\Delta\chi$ [Mpc]	$d\lambda/d\chi$	$dv/d\chi$	$\lvert b \rvert$
2.0	3647	0.88	782	0.82	67	0.12
2.5	4255	0.80	740	1.01	71	0.18
3.0	4863	0.70	701	1.22	75	0.27
3.5	5471	0.58	666	1.45	79	0.37
4.0	6078	0.44	635	1.69	83	0.55

2 SENSITIVITY TO 3D FLUX POWER SPECTRUM

2.1 Covariance

Let us take a quasar survey that provides N Lyα forest spectra at redshift z. It provides these Lyα forest spectra towards the locations $\boldsymbol{x}_{\perp,n}$ on the sky, where $n \in \{1, ..., N\}$. We treat the line-of-sight direction as continuous for convenience such that the survey measures the Lyα forest in a spatial window given by

$$W(\boldsymbol{x}) = \bar{n}^{-1} \sum_{n=1}^{N} \delta^D \left(\boldsymbol{x}_\perp - \boldsymbol{x}_{\perp,n}\right), \tag{1}$$

where $\bar{n} \equiv N/\mathcal{A}$, $\mathcal{A}$ is the survey area on the sky, and $\boldsymbol{x} = (x_\parallel, \boldsymbol{x}_\perp)$. We use an analogous convention for the Fourier wavevector, $\boldsymbol{k}$. Thus, the survey measures $\delta(\boldsymbol{x}) = \delta_{\rm F}(\boldsymbol{x})\, W(\boldsymbol{x})$, where $\delta_{\rm F}(\boldsymbol{x}) = F/\langle F \rangle - 1$ is the overdensity in the transmission fraction F, where $\langle F \rangle$ is the mean transmission. Switching to the Fourier basis, the covariance of two δ-modes is

$$\begin{aligned} \langle \tilde{\delta}(\boldsymbol{k}) \tilde{\delta}(\boldsymbol{k}')^* \rangle &= \bar{n}^{-2}\, (2\pi)\, \delta^D(k_\parallel - k_\parallel') \qquad (2)\\ &\times \sum_{n=1}^{N} \sum_{m=1}^{N} \tilde{w}_n(k_\parallel) \tilde{w}_m(k_\parallel')^* e^{-i(\boldsymbol{k}_\perp \cdot \boldsymbol{x}_{\perp,n} - \boldsymbol{k}'_\perp \cdot \boldsymbol{x}_{\perp,m})} \\ &\times \left[P_{{\rm N},n}\, \delta_n^m + \int \frac{d^2\boldsymbol{k}_\perp''}{(2\pi)^2} e^{i\boldsymbol{k}_\perp'' \cdot (\boldsymbol{x}_{\perp,n} - \boldsymbol{x}_{\perp,m})} P_{\rm F}(k_\parallel, \boldsymbol{k}_\perp'') \right], \end{aligned}$$

where $\langle ... \rangle$ represents an ensemble average, $P_{{\rm N},n}$ is the 1D noise power for skewer n (which we assume does not correlate with other sightlines), δ_n^m is the Kronecker delta, tildes signify the Fourier basis, and $\tilde{w}_n(k_\parallel)$ is the weight given to the mode for skewer n and normalized such that $N^{-1} \sum_n \tilde{w}_n(k_\parallel) = 1$. We will often keep the $k_\parallel$ dependence of $\tilde{w}_n$ implicit for notational simplicity. This weighting scheme assumes that the combined weight of pixels n and m factorizes. While not fully general, with the proper choice of w_n, this scheme is close to optimal (Section 3). Lastly, equation (2) approximates the modes as continuous in the transverse direction, as would occur in the limit $\mathcal{A} \to \infty$.

The noise power that appears in equation (2) can be related to the $[S/N]_X$ on the continuum in a pixel in a resolution element of size X:

$$\begin{aligned} P_{{\rm N},n} &= \langle F \rangle^{-2} \, [S/N]_{\Delta x}^{-2} \left(\frac{\Delta x}{1\,{\rm Mpc}} \right), & (3)\\ &= 0.8 \, \langle F \rangle^{-2} \, [S/N]_{\Delta \lambda}^{-2} \left(\frac{\Delta \lambda}{1 \text{ Å}} \right) \left(\frac{1+z}{4} \right)^{-3/2}, & (4) \end{aligned}$$

where $\Delta\lambda$ is the pixel size in wavelength, and we have assumed that the noise is white. We will henceforth discuss the noise in terms of the S/N in 1 Å pixels. However, for the tens of Mpc modes of interest for 3D Lyα forest analyses, resolving 1 Å may not be required. A resolution of tens of Å suffices to capture the scales where 3D correlations can be detected for upcoming measurements (McDonald & Eisenstein 2007). It is simply a convenience that leads us to quote the S/N per 1 Å pixel.

Equation (2) is the covariance of the Fourier-space flux overdensity field, $\tilde{\delta}_{\rm F}$, convolved with the survey window, $\widetilde{W}$. In the limit of large N, $\widetilde{W}$ becomes a δ-function such that $\tilde{\delta} \approx \tilde{\delta}_{\rm F}$. The covariance of two modes in the flux field can be approximated in this limit as diagonal with entries[2]

$$P_{\rm tot} \equiv \langle |\tilde{\delta}_k|^2 \rangle = P_{\rm F}(\boldsymbol{k}) + \bar{n}^{-1} \left[P_{\rm N} + \overline{\tilde{w}^2}\, P_{\rm los}(k_\parallel) \right], \tag{5}$$

where $\overline{X}$ indicates the average of quantity X_n over N skewers,

$$P_{\rm los}(k_\parallel) \equiv \int \frac{d^2\boldsymbol{k}_\perp}{(2\pi)^2}\, P_{\rm F}(k_\parallel, \boldsymbol{k}_\perp), \tag{6}$$

and

$$P_{\rm N} \equiv \frac{1}{N} \sum_n^N |\tilde{w}_n|^2 P_{{\rm N},n}. \tag{7}$$

The function $P_{\rm los}$ is just the line-of-sight power spectrum that is typically measured in cosmological studies of the Lyα forest (e.g., McDonald et al. 2005). As in McDonald & Eisenstein (2007), we will refer to this term in equation (6) as the "aliasing" term. Equation (5) is derived by noting that the sum over N skewers is a monte-carlo evaluation of the integral, and the monte-carlo integral of a plane wave leads to a δ-function to precision of $N^{-1/2}$. Replacing these summations with δ-functions results in a two δ-function term from the components that have $m \neq n$ (that yields $P_{\rm F}$) and a one δ-function term from $m = n$ (that yields $P_{\rm tot} - P_{\rm F}$). Equation (5) agrees with equation (13) in McDonald & Eisenstein (2007) and with equation (A6) in White et al. (2010).[3]

Henceforth we ignore correlations between flux pixels in a single line-of-sight, which contain a small fraction of the total information. This omission is equivalent to enforcing the condition that $m \neq n$ in the summation in equation (2). We define the estimator $\widehat{P}_{\rm F} \equiv |\tilde{\delta}_{\boldsymbol{k}}|^2$ for this case. Dropping terms suppressed by higher powers of N, the ensemble average of $\widehat{P}_{\rm F}$ is

$$\langle \widehat{P}_{\rm F}(\boldsymbol{k}) \rangle = P_{\rm F}(\boldsymbol{k}). \tag{8}$$

Thus, omitting line-of-sight correlations results in an unbiased estimator at lowest order. While $\widehat{P}_{\rm F}$ is biased to fractional order $N^{-1/2}$, the noise in this estimate can be corrected because the $\boldsymbol{x}_{\perp,n}$ are known as shown in Section 3. An additional advantage of omitting line-of-sight correlations is that the contamination from continuum fluctuations and damping wings primarily induce this type of correlation (Section 5).

The covariance of $\widehat{P}_{\rm F}(\boldsymbol{k})$ is

$$\begin{aligned}
{\rm cov}[\widehat{P}_{\rm F}(\boldsymbol{k}),\ \widehat{P}_{\rm F}(\boldsymbol{k}')]_{k_\parallel} &= \mathcal{A}^2\, \bar{n}^{-4} \int \frac{d^2\boldsymbol{k}_\perp^{(1)}\, d^2\boldsymbol{k}_\perp^{(2)}}{(2\pi)^4} \sum_{n,m,k,l}^{n\neq m, k\neq l} \\
\times \quad & e^{i\boldsymbol{x}_n\cdot(\boldsymbol{k}_\perp - \boldsymbol{k}_\perp^{(1)})} e^{-i\boldsymbol{x}_m\cdot(\boldsymbol{k}_\perp - \boldsymbol{k}_\perp^{(2)})} \Big[e^{i\boldsymbol{x}_k\cdot(\boldsymbol{k}'_\perp - \boldsymbol{k}_\perp^{(2)})} \\
& \times e^{-i\boldsymbol{x}_l\cdot(\boldsymbol{k}'_\perp - \boldsymbol{k}_\perp^{(1)})} + e^{i\boldsymbol{x}_k\cdot(\boldsymbol{k}'_\perp + \boldsymbol{k}_\perp^{(2)})} e^{-i\boldsymbol{x}_l\cdot(\boldsymbol{k}'_\perp + \boldsymbol{k}_\perp^{(1)})} \Big] \\
\times \quad & \tilde{w}_n\, \tilde{w}_l^* \left[P_{\rm F}(k_\parallel, \boldsymbol{k}_\perp^{(1)}) + P_{{\rm N},n}\, \delta_n^l\, (2\pi)^2\, \delta^D(\boldsymbol{k}_\perp^{(1)}) \right] \\
\times \quad & \tilde{w}_k\, \tilde{w}_m^* \left[P_{\rm F}(k_\parallel, \boldsymbol{k}_\perp^{(2)}) + P_{{\rm N},m}\, \delta_m^k\, (2\pi)^2\, \delta^D(\boldsymbol{k}_\perp^{(2)}) \right],
\end{aligned} \tag{9}$$

where all 3D wavevectors that appear are evaluated at $k_\parallel$, and we have assumed Gaussianity and, thus, dropped the connected 4th moment. While the small-scale modes in the forest are far from being Gaussian, Gaussianity is likely to be a decent approximation at the > 10 Mpc wavelength modes of interest for 3D analyses. With a bit of algebra analogous to how equation (5) was derived, this equation simplifies to

$$\begin{aligned}
{\rm cov}[\widehat{P}_{\rm F}(\boldsymbol{k}), \widehat{P}_{\rm F}(\boldsymbol{k}')]_{k_\parallel} &= 2P_{\rm tot}^2 \delta_{\boldsymbol{k}}^{\boldsymbol{k}'} + \tfrac{4}{N}\, \overline{\tilde{w}^2}\, P_{\rm F}(\boldsymbol{k}) P_{\rm F}(\boldsymbol{k}') \\
+ \tfrac{1}{N}\, \overline{\tilde{w}^2}^2\, \bar{n}^{-1} & \int \tfrac{d^2\boldsymbol{k}_\perp^{(1)}}{(2\pi)^2}\, P_{\rm F}(k_{\rm ll}, \boldsymbol{k}_\perp^{(1)}) P_{\rm F}(k_{\rm ll}, \boldsymbol{k}_\perp^{(1)} - \boldsymbol{k}_\perp - \boldsymbol{k}'_\perp) \\
+ \tfrac{1}{N}\, \overline{\tilde{w}^2}^2\, \bar{n}^{-1} & \int \tfrac{d^2\boldsymbol{k}_\perp^{(1)}}{(2\pi)^2}\, P_{\rm F}(k_{\rm ll}, \boldsymbol{k}_\perp^{(1)}) P_{\rm F}(k_{\rm ll}, \boldsymbol{k}_\perp^{(1)} + \boldsymbol{k}_\perp - \boldsymbol{k}'_\perp).
\end{aligned} \tag{10}$$

For the wavevectors and quasar number densities of interest, the off-diagonal terms in equation (10) are unimportant. The subtlety here is that, unlike the diagonal terms, these terms do not average down linearly with the number of independent modes when binning the estimates for $P_{\rm F}$ in a shell in k-space. This is analogous to the well-known behavior of the 4-point function in surveys of large-scale structure (Meiksin & White 1999), which has become known as "beat-coupling" (Hamilton et al. 2006; Rimes & Hamilton 2006). The most important off-diagonal contribution comes from $4\,N^{-1}\, P_{\rm F}(\boldsymbol{k}) P_{\rm F}(\boldsymbol{k}')$. At fixed $k_\parallel$, this becomes comparable to the corresponding diagonal term when the number of binned pixels in the shell is comparable to N. This criterion is satisfied for larger wavevectors than $k_\perp = 0.06\, [\bar{n}/10^{-3}\, {\rm Mpc}^{-2}]^{1/2} \times \sqrt{k_\perp/\Delta k_\perp}\, {\rm Mpc}^{-1}$, where $\Delta k_\perp$ is the size of the bin in the $k_\perp$ direction. However, for the wavevectors and number densities of relevance, we find that in practice this term is always subdominant to the aliasing term, $\bar{n}^{-1}\, P_{\rm F}$, in $P_{\rm tot}$. The inclusion of these terms is however formally important in that they serve to "cap" the total information content of a survey, resolving an issue raised in McDonald & Eisenstein (2007). The off-diagonal elements are a more important consideration when line-of-sight correlations are included.

Quasar clustering was not accounted for in the previous expressions. While clustering does not bias $\widehat{P}_F$ when line-of-sight correlations are omitted (at least at quadratic order in the density), it does increase the variance of the estimate. Clustering results in the term $\bar{n}^{-1}\, (P_{\rm N} + P_{\rm los})$ in $P_{\rm tot}$ gaining the factor $1 + C_q(\boldsymbol{k}_\perp)\, \bar{n}$, where $C_q(\boldsymbol{k}_\perp)$ is the angular power spectrum of the sources that contribute at that redshift.

---

[2] The largest of the off-diagonal terms is $P_{\rm F}\, {\rm VAR}[\overline{\tilde{w}^2}/N]$, where the "VAR" function represents a number whose probability distribution function has a variance given by the argument. Its exact value is determined by the $\boldsymbol{x}_{\perp,n}$. Even though there are N^2 possible correlations, the error from neglecting the off-diagonal terms scales as $N^{-1/2}$ rather than N^{-1} because there are only N quasar locations.

[3] As a consistency check, it is simple to show that $\langle \chi^2(k_\parallel) \rangle \equiv \sum P_{\rm F}/P_{\rm tot}$ is $\leqslant N$, where the sum is over all wavevectors with the same $k_\parallel$. This inequality demonstrates that the amount of information that can be extracted from $\tilde{\delta}_F$ (as measured by the significance at which $\tilde{\delta}_F$ can be detected) is less than or equal to the number of Lyα forest pixels in a survey.

Table 2. Sloan Digital Sky Survey constraints on $P_{\rm los}$ in units of Mpc from McDonald et al. (2005).

$k_\parallel$	$z = 2.2$	$z = 2.6$	$z = 3.0$	$z = 3.6$	$z = 4.0$
0.15	0.27(2)	0.43(2)	0.58(4)	1.05(10)	1.47(22)
0.20	0.26(1)	0.38(2)	0.58(3)	0.86(6)	1.08(13)
0.30	0.18(1)	0.30(1)	0.44(2)	0.81(5)	0.85(10)
0.50	0.15(1)	0.24(1)	0.35(1)	0.59(3)	0.81(07)

Table 3. Value of $\nu_n \equiv (1+P_{{\rm N},n}/P_{\rm los})^{-1}$ (c.f., eqn. 13) at several redshifts as a function of the S/N on the continuum in a 1 Å pixel. This calculation uses the estimate of McDonald et al. (2005) at $k = 0.0014$ s km^{-1} for $P_{\rm los}$.

$[S/N]_{1A}$	$z = 2.2$	$z = 2.6$	$z = 3.0$	$z = 3.6$	$z = 4.0$
0.50	0.04	0.06	0.08	0.10	0.11
1.00	0.14	0.20	0.27	0.31	0.33
2.00	0.40	0.50	0.59	0.64	0.66
5.00	0.81	0.86	0.90	0.92	0.93
10.00	0.94	0.96	0.97	0.98	0.98

This correction is significant on scales where the clustering power is comparable or larger than the shot power. Assuming a quasar bias of 3 at $z = 3$ and that each Lyα spectrum yields 500 Mpc of absorption, the two become comparable for $\bar{n} \approx 10^{-2}$ Mpc^{-2} (and then only at $k_\perp \approx 0.02$ Mpc^{-1}, the horizon scale at matter-radiation equality). This $\bar{n}$ is an order of magnitude higher than what upcoming surveys will achieve.

2.2 Optimal Weights

The weight that maximizes the signal-to-noise ratio, $P_{\rm F}^2/{\rm var}[\widehat{P}_{\rm F}]$, is

$$\tilde{w}_n(k_\parallel) = \mathcal{B}(k_\parallel)\ \left(P_{\rm los}(k_\parallel) + P_{{\rm N},n}\right)^{-1}, \qquad (11)$$

where $\mathcal{B}(k_\parallel)$ is fixed by our normalization convention and absorbs all factors that do not depend on n, and this maximization neglects the subdominant off-diagonal covariances so that ${\rm var}[\widehat{P}_{\rm F}] = 2\,P_{\rm tot}^2$. (We show in Section 2.4 that a simple generalization of these weights can be easily applied to real data.) This choice of weights results in $P_{\rm tot}$ becoming

$$P_{\rm tot}(\boldsymbol{k}) = P_{\rm F}(\boldsymbol{k}) + \bar{n}_{\rm eff}^{-1}\,P_{\rm los}(k_\parallel), \qquad (12)$$

where

$$\bar{n}_{\rm eff} \equiv \frac{1}{\mathcal{A}}\sum_{n=1}^{N}\nu_n, \quad \nu_n \equiv \frac{P_{\rm los}(k_\parallel)}{P_{\rm los}(k_\parallel) + P_{{\rm N},n}}. \qquad (13)$$

Table 4. $\bar{n}_{\rm eff}$ for a quasar survey in units of 10^{-3} Mpc^{-2} assuming that $m_{\rm AB}^{\rm lim} = m_{\rm AB}^{1A}$ and $k = 0.1$ Mpc^{-1}. These numbers are calculated using the B-band luminosity function from Hopkins et al. (2006). The second number in select entries is this but also using galaxies in addition to quasars. The bottom row provides the factor that converts $\bar{n}_{\rm eff}$ in each respective column to units of deg^{-2}.

$m_{\rm AB}^{1A}$	$z = 2$	$z = 2.5$	$z = 3$	$z = 4$	$z = 5$
21	0.26	0.16	0.07	0.01	0.00
22	0.82	0.54	0.27	0.04	0.00
23	2.0	1.4	0.78	0.11	0.01
24	3.8/5	2.8	1.7	0.32	0.04
25	6.4/18	4.8/9	3.2/5	0.77	0.10
26	9.8/66	7.4/47	5.1/28	1.5/4	0.24
	8.4	11	13	16	19

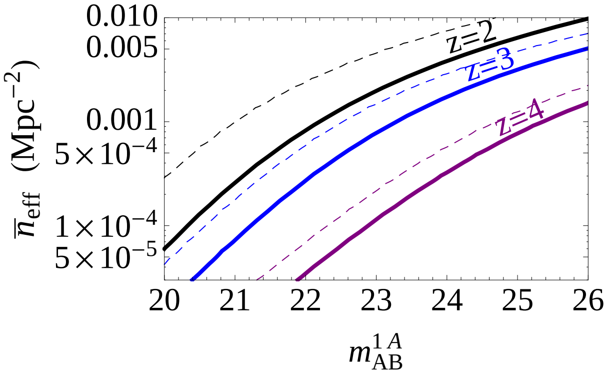

Figure 1. Solid curves show the effective number density of quasars contributing Lyα forest spectra at redshift z (eqn. 13) as a function of the B-band AB magnitude that has $[S/N]_{1A} = 1$, $m_{\rm AB}^{1A}$. We have also assumed that $m_{\rm AB}^{1A}$ is equal to the limiting magnitude of the survey. The dashed curves are the actual number of quasars brighter than $m_{\rm AB}$. These curves are calculated for $k_\parallel = 0.1$ Mpc^{-1} using the Hopkins et al. (2006) luminosity function and assuming that $S/N \propto$ flux.

Thus, a single number, $\bar{n}_{\rm eff}$, characterizes the sensitivity of a Lyα forest survey to $P_{\rm F}$, and ν_n is a measure of the importance of each quasar on a scale of 0 to 1. While $\bar{n}_{\rm eff}$ depends on $k_\parallel$, in practice this dependence is likely to be weak because $P_{\rm los}(k_\parallel)$ has roughly a white noise power spectrum at relevant wavevectors and $\bar{n}_{\rm eff}$ also depends relatively weakly on $P_{\rm los}(k_\parallel)$. The constancy of $P_{\rm los}$ at $k \leqslant 0.5$ Mpc^{-1} is quantified in Table 2, which tabulates measurements of $P_{\rm los}$ from McDonald et al. (2005) at several $k_\parallel$ and redshifts. Constancy should be an even better approximation at smaller $k_\parallel$ than is tabulated. For a BOSS-like Lyα forest survey at $z = 2.5$, a factor of 2 smaller $P_{\rm los}(k_\parallel)$ from its small-$k_\parallel$ asymptote results in only a factor of 1.4 decrease in $\bar{n}_{\rm eff}$. The decrease is even smaller for a deeper survey.

The gains in sensitivity to $P_{\rm F}$ are meager from improving the S/N on a quasar once $P_{{\rm N},n} < P_{\rm los}$, which corresponds roughly to $[S/N]_{1A} \approx 2$. Table 3 quantifies this statement by giving ν_n as a function of the S/N at 1 Å, using the estimates of $P_{\rm los}$ at the smallest k quoted in McDonald et al. (2005) of 0.0014 s km^{-1}. However, these numbers should be applicable over a wide range of k owing to the form of $P_{\rm los}$. Even though the amount of power in the forest increases with redshift, the S/N requirements at fixed source flux remain nearly constant with increasing redshift (becoming slightly less stringent).

Figure 1 shows $\bar{n}_{\rm eff}$ at three redshifts as a function of the B-band magnitude at which a survey obtains $[S/N]_{1\,A} = 1$. For this and subsequent calculations, we assume that $P_{\rm F}$ has the simple form

$$P_{\rm F}(\boldsymbol{k}, z) = b(z)^2\,(1 + g\,\mu^2)^2\,P_\delta^{\rm lin}(k, z)\,\exp[-k_\parallel^2/k_{\rm D}^2], \quad (14)$$

where $P_\delta^{\rm lin}$ is the linear-theory density power spectrum and $k_{\rm D}^2 \equiv m_p/[k_{\rm B}T]$. We assume $T = 20,000$ K isothermal gas such that $k_{\rm D} = 0.08$ s^{-1} km, and $b(z)$ is calibrated so that $P_{\rm los}$ matches the McDonald et al. (2005) measurements at 0.0014 s^{-1} km (see Table 1). The factor $(1 + g\,\mu^2)^2$ owes to peculiar velocities, where $\mu = \hat{\boldsymbol{n}} \cdot \hat{\boldsymbol{k}}$ and $\hat{\boldsymbol{n}}$ is the unit vector along the line-of-sight. We set $g = 1$, a choice motivated by the results of Slosar et al. (2009) and McQuinn et al. (2010), but caution that there is at present theoretical uncertainty in this choice at the 50% level (McQuinn et al. 2010).

The solid curves in Figure 1 are calculated using the Hopkins et al. (2006) B-band luminosity function and assuming that each quasar contributes ≈ 500 Mpc of Lyα absorption (specifically the forest between 1041 and 1185 Å in each sightline). Thus, these curves are the effective number density of quasars projected over this distance, which can be thought of as the number of quasars that contribute Lyα forest spectra at the stated redshift. These curves also take the survey limiting magnitude $m_{\rm AB}^{\rm lim}$ to be equal to the magnitude where quasars have $[S/N]_{1A} = 1$, $m_{\rm AB}^{1A}$, a choice we discuss below, and assume that $S/N \propto$ flux. This scaling corresponds to the sensitivity being sky- or dark current-limited, a choice motivated by the fact that this scaling matters most for the faintest objects that are observed. However, the value of $\bar{n}_{\rm eff}$ is not significantly changed if rather we assume $S/N \propto {\rm flux}^{1/2}$ as would occur if the observations were photon-limited. The solid curves from top to bottom are $\bar{n}_{\rm eff}$ for $z = 2,\ 3$, and 4. These curves illustrate that it will be challenging to (1) measure the 3D Lyα forest at $z > 3$ because of the falloff in the total abundance of quasars or (2) obtain $\bar{n}_{\rm eff} \gg 10^{-3}$ Mpc^{-2} at any redshift because of the shallow faint-end slope of the luminosity function. Table 4 tabulates $\bar{n}_{\rm eff}$ as various redshifts and limiting magnitudes, and it also gives the conversion from Mpc^{-2} to deg^{-2} units.

The dashed curves in Figure 1 show the total number of quasars brighter than $m_{\rm AB}$. The effective number density at $m_{\rm AB}^{1A}$ is always a factor of a few smaller than the total number of quasars, with the difference decreasing with redshift. See Appendix A for an analytic understanding of this suppression for power-law luminosity functions.

BOSS will achieve a B-band magnitude limit of $m_{\rm AB}^{\rm lim} \approx 22$ at $z = 2$ ($m_{\rm AB}^{\rm lim} \approx 21$ at $z = 3$) and also achieve $m_{\rm AB}^{1A} \approx 22$. These numbers yield $\bar{n}_{\rm eff} = (0.3 - 0.8) \times 10^{-3}$ Mpc^{-2} [$4 - 7$ deg^{-2}] at $z = 2 - 3$. BigBOSS aims to achieve one magnitude fainter than BOSS, which results in $\bar{n}_{\rm eff} = (1 - 2) \times 10^{-3}$ Mpc^{-2} at $z = 2 - 3$. For these $\bar{n}_{\rm eff}$ and at $k \gtrsim 0.1$ Mpc^{-1}, S/N scales linearly with $\bar{n}_{\rm eff}$ such that BigBOSS will be a few times more sensitive than BOSS.

Figure 2 is a contour plot of $\bar{n}_{\rm eff}$ as a function of $m_{\rm AB}^{\rm lim}$ and $m_{\rm AB}^{1A}$. The solid contours are calculated assuming $k_\parallel = 0.1$ Mpc^{-1} and are labeled in units of 10^{-3} Mpc^{-2}. These curves illustrate that $\bar{n}_{\rm eff}$ is maximized roughly when $m_{\rm AB}^{\rm lim} \approx m_{\rm AB}^{1A} - 0.5$. Little is gained by observing fainter quasars than $m_{\rm AB}^{1A} - 0.5$ or by integrating longer on the same quasars once $m_{\rm AB}^{1A} \approx m_{\rm AB}^{\rm lim} + 0.5$. The dashed contours are the same but at $k_\parallel = 0.5$ Mpc^{-1} (at which $P_{\rm los}$ is a factor of 0.6 smaller in our model), demonstrating that the $k_\parallel$ dependence of these conclusions is weak.

How much is gained by weighting by w_n relative to a uniform weighting scheme? Let us take a luminosity with power-law index -2, a form discussed in Appendix A. We assume that $S/N \propto$ flux, $P_{\rm N,lim}/P_{\rm los} = 1$ (equivalent to $[S/N]_{1A} \approx 2$; Table 3), where $P_{\rm N,lim}$ is the noise power spectrum for quasars at the survey limiting magnitude. This case results in only a 4% improvement in the value of $P_{\rm tot} - P_{\rm F}$ relative to a simple uniform weighting. For $P_{\rm N,lim}/P_{\rm los} = 0.5\,(0.3)$, the improvement is 40% (60%). Thus, weighting only offers a significant improvement when a large fraction of the quasars have $P_{\rm N,n}/P_{\rm los} < 1$.

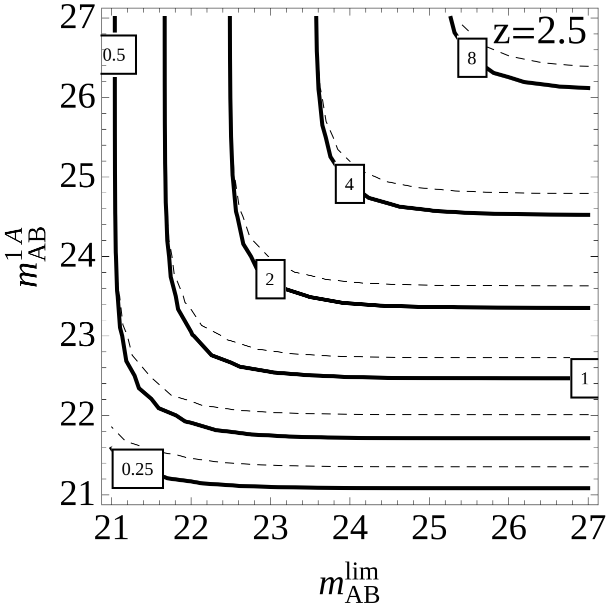

Figure 2. Impact of varying the limiting magnitude of the survey, $m_{\rm AB}^{\rm lim}$, along with the magnitude at which $[S/N]_{1A} = 1$, $m_{\rm AB}^{1A}$. The solid curves are contours of constant $\bar{n}_{\rm eff}$ evaluated at $k_\parallel = 0.1$ Mpc^{-1} and $z = 2.5$, with the labels in units of 10^{-3} Mpc^{-2}. The corresponding dashed curves are the same but evaluated at $k_\parallel = 0.5$ Mpc^{-1}.

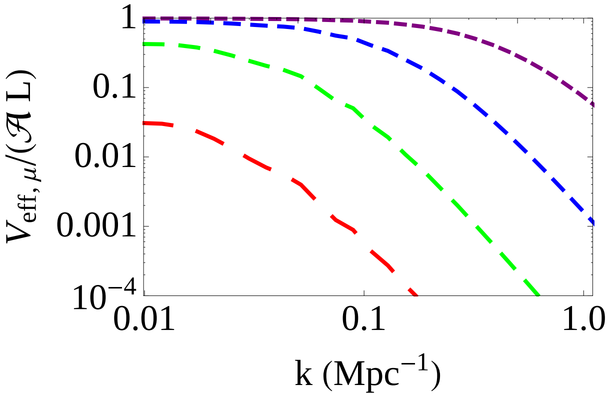

Figure 3. 3D Lyα forest effective volume divided by survey volume at $|\boldsymbol{k}|$, $z = 3$, and for $\bar{n}_{\rm eff} = 10^{-1}$, 10^{-2}, 10^{-3}, and 10^{-4} Mpc^{-2} in order of increasing dash length (decreasing amplitude). The plotted quantity is equal to $(2\,P_{{\rm F},\mu}/\delta P_{{\rm F},\mu})^2$ and is essentially independent of redshift at the plotted scales.

2.3 Effective Volume

One can define an effective volume for Lyα forest surveys analogous to the effective volume in galaxy surveys

(Feldman et al. 1994). In particular, in the galaxy survey case the effective volume is

$$V_{\rm eff,g}(\boldsymbol{k}) \equiv V_{\rm g} \left(\frac{P_{\rm g}(\boldsymbol{k})}{P_{\rm g}(\boldsymbol{k}) + \bar{n}_{\rm g,3D}^{-1}} \right)^2, \quad (15)$$

which becomes in the Lyα forest survey case

$$V_{\rm eff}(\boldsymbol{k}) \equiv \mathcal{A}\, L \left(\frac{P_{\rm F}(\boldsymbol{k})}{P_{\rm F}(\boldsymbol{k}) + \bar{n}_{\rm eff}^{-1} P_{\rm los}(k_\parallel)} \right)^2. \quad (16)$$

In equation (15) and (16), $V_{\rm g}$ is the actual volume of the galaxy survey, $P_{\rm g}$ is the galaxy power spectrum, $\bar{n}_{\rm g,3D}$ is the number density of galaxies, and L is the line-of-sight dimension of the Lyα survey. The major difference between these two effective volumes is just that in the Lyα case the shot noise is in the plane of the sky and this term is modulated by the line-of-sight power.

Figure 3 plots a generalization of the the effective volume, $V_{{\rm eff},\mu}(k)/[\mathcal{A}\, L]$, for $\bar{n}_{\rm eff} = 10^{-1}$, 10^{-2}, 10^{-3}, and 10^{-4} Mpc^{-2}, in order of increasing dash length. We define $V_{{\rm eff},\mu}(k)$ to be $V_{\rm eff}(\boldsymbol{k})$ but averaged over a shell in k-space such that $V_{{\rm eff},\mu}(k) \equiv 2\, \mathcal{A} L\, P_{{\rm F},\mu}^2 / \delta P_{{\rm F},\mu}^2$, where $P_{{\rm F},\mu} \equiv P_F/(1+g\mu^2)^2$,

$$\delta P_{{\rm F},\mu} = \left(\int_0^1 d\mu \, (1+g\mu^2)^4 \, {\rm var}[\widehat{P}_{\rm F}(k,\mu)]^{-1} \right)^{-1/2}, \quad (17)$$

and ${\rm var}[\widehat{P}_{\rm F}(\boldsymbol{k})] \equiv {\rm cov}[\widehat{P}_{\rm F}(\boldsymbol{k}), \widehat{P}_{\rm F}(\boldsymbol{k})]$. Here, $\delta P_{{\rm F},\mu}$ has been defined so that $\delta P_{{\rm F},\mu}/\mathcal{N}_k^{1/2}$ is the precision at which $P_{{\rm F},\mu}$ can be measured in a k-space shell with $\mathcal{N}_k$ independent modes.

The value of $V_{{\rm eff},\mu}(k)/[\mathcal{A}\, L]$ is near unity when a survey is sample variance-limited. For a BOSS-like quasar survey in which $\bar{n}_{\rm eff} = 10^{-3}$ Mpc^{-2} at $2 < z < 3$, this quantity reaches a maximum of 0.5 at the smallest wavevectors and falls off rapidly at $k \gtrsim 0.1$ Mpc^{-1}. Even $\bar{n}_{\rm eff} = 10^{-2}$ Mpc^{-2} – roughly the maximum density that can be achieve in a quasar survey – is not sample variance-limited at $k \gtrsim 0.1$ Mpc^{-1}. Figure 3 also emphasizes that $\bar{n}_{\rm eff} \approx 10^{-3}$ Mpc^{-2} marks the critical number density at which a survey becomes sample variance-limited at $k \lesssim 0.1$ Mpc^{-1}, such that the gains from a 3D analysis become large.

The plotted quantity in Figure 3 can also be related to the S/N of quasar survey to $P_{{\rm F},\mu}$ in a k-space shell of width Δk. This shell contains $\mathcal{N}_k \equiv k^2 \Delta k\, \mathcal{A}\, L/(2\pi^2)$ independent samples, so that the S/N ratio at which $P_{{\rm F},\mu}$ can be measured in this shell equals

$$\begin{aligned} P_{{\rm F},\mu}/\sigma_{{\rm F},\mu} \;&=\; 60 \left(\frac{V_{{\rm eff},\mu}/\mathcal{A}\, L}{10^{-2}} \right)^{1/2} \left(\frac{L}{1\,{\rm Gpc}} \frac{\mathcal{A}}{10^3\,{\rm deg}^2} \right)^{1/2} \\ &\times \left(\frac{1+z}{4} \right)^{1.2} \left(\frac{k^3\,[\Delta k/k]}{10^{-3}\,{\rm Mpc}^{-3}} \right)^{1/2}, \end{aligned} \quad (18)$$

where $\sigma_{{\rm F},\mu} \equiv \delta P_{{\rm F},\mu}(k)/\mathcal{N}_k^{1/2}$ and the redshift dependence is approximate.

2.4 Estimating Correlations on Real Data

In practice, it will be easier to measure 3D Lyα correlations in configuration space via the 3D correlation function. In configuration space, the Fourier-space weights proposed in Section 2.2 become the line-of-sight convolution of $w_n(r)$ – the Fourier Transform of $\tilde{w}_n(k_\parallel)$ – with the flux field of sightline n. Conveniently, $w_n(r)$ will be relatively localized in real-space because $\tilde{w}_n(k_\parallel) \approx \tilde{w}_n(0)$ for $k_\parallel \lesssim 0.5$ Mpc^{-1}. However, for estimates of the correlation function, a simpler weighting for each sightline of $\tilde{w}_n(0)$ is likely sufficient. In fact, this weighting is identical to the full weighting to the extent that $P_{\rm los}$ and the $P_{\rm N,n}$ are white. We find that ${\rm var}[\widehat{P}_{\rm F}]$ is not significantly increased with these simpler weights.

Real data will be more complicated than the idealized case of uniform noise that we have considered thus far. For example, there will be high-noise regions in the spectra owing to sky lines, and there may also be holes in the data where damped systems have been excised. We have expressed our weights in terms of the noise and the Lyα forest power spectra. However, the suggested weights are equivalent to simply weighting each sightline by $(1 + \sigma_{\rm N}^2/\sigma_{\rm los}^2)^{-1}$, where $\sigma_{\rm N}^2$ and $\sigma_{\rm los}^2$ are respectively the variance of the noise and the Lyα forest, smoothed on a large enough scale that $P_{\rm los}$ is white. This scale corresponds to $\gtrsim 10$ Mpc, and we find that the exact choice of the smoothing scale weakly affects ${\rm var}[\widehat{P}_{\rm F}]$. Such weights can more readily be applied to real data.

3 QUADRATIC ESTIMATOR

This section derives the minimum variance estimator that is quadratic in $\delta_{\rm F}$ and compares it to the estimator derived in the Section 2. To proceed, we decompose the covariance of the flux overdensity at $k_\parallel$ into a component that depends on the parameter we aim to measure, $P_{\rm F}(\boldsymbol{k})$, and one that does not. Namely,

$$\left\langle \tilde{\delta}_{{\rm F},n}\, \tilde{\delta}_{{\rm F},m} \right\rangle = C_{nm} + P_{\rm F}(\boldsymbol{k})\, X_{nm}, \quad (19)$$

where $\tilde{\delta}_{{\rm F},n}$ is the Fourier transform of the flux along sightline n and we consider a single $k_\parallel$,

$$\begin{aligned} C_{nm} \;&=\; P(k_\parallel, \boldsymbol{r}_{nm}) + P_{{\rm N},n}\, \delta_n^m, & (20)\\ X_{nm} \;&=\; \left(\frac{\Delta k_\perp}{2\pi} \right)^2 \exp[i \boldsymbol{k}_\perp \cdot \boldsymbol{r}_{nm}], & (21)\\ P(k_\parallel, \boldsymbol{r}_{nm}) \;&\approx\; \int \frac{dk_\perp}{2\pi}\, k_\perp\, P_{\rm F}(k_\parallel, k_\perp)\, J_0(k_\perp\, r_{nm}). & (22) \end{aligned}$$

Here, $r_{nm} \equiv |\boldsymbol{r}_{nm}|$ is the transverse separation between sightlines n and m, $\Delta k_\perp$ is the width of the Fourier space pixel in which we estimate $P_{\rm F}(\boldsymbol{k})$, and equation (22) assumes that $P_{\rm F}(\boldsymbol{k})\, X_{nm} \ll P(k_\parallel, \boldsymbol{r}_{nm})$.

The minimum variance estimator for $P_{\rm F}(\boldsymbol{k})$ that is quadratic in $\tilde{\delta}_{{\rm F},n}$ and unbiased is given by

$$\widehat{P}_{\rm F}^{\rm QE} = \sum_{nm} Q_{nm}\, \tilde{\delta}_{{\rm F},n} \tilde{\delta}_{{\rm F},m} - {\rm tr}[\boldsymbol{Q}\, \boldsymbol{C}], \quad (23)$$

where

$$\boldsymbol{Q} = \left({\rm tr}\left[\boldsymbol{C}^{-1} \boldsymbol{X}\, \boldsymbol{C}^{-1} \boldsymbol{X} \right] \right)^{-1} \boldsymbol{C}^{-1} \boldsymbol{X}\, \boldsymbol{C}^{-1}. \quad (24)$$

Equation (24) is derived by minimizing the variance of $\widehat{P}_{\rm F}^{\rm QE}$ assuming $\delta_{\rm F}$ is Gaussian so that

$${\rm var}\left[\widehat{P}_{\rm F}^{\rm QE} \right] = 2\, {\rm tr}\left[\boldsymbol{C}\, \boldsymbol{Q}\, \boldsymbol{C}\, \boldsymbol{Q} \right], \quad (25)$$

and subject to the condition that it is unbiased so that ${\rm tr}[\boldsymbol{Q}\, \boldsymbol{X}] = 1$. Note that equation (25) does not include sample variance, which would contribute the term $2\, P_{\rm F}^2$.

Our aim is to understand how the weights suggested in the previous section relate to the above minimum variance quadratic estimator. To proceed, we set the off diagonal terms in C_{nm} to zero (i.e., $P(k_\parallel, \boldsymbol{r}_{nm}) \to 0$ for $\boldsymbol{r}_{nm} \neq 0$). The estimator is still unbiased with this approximation. The motivation for this approximation is that the typical separation of quasars in such surveys is $r_\perp \gtrsim 10$ Mpc. At these $r_\perp$ and at relevant $k_\parallel$, $P(k_\parallel, \boldsymbol{r}_\perp)$ is down by a factor of $\gtrsim 4$ from its $\boldsymbol{r}_\perp = 0$ value. In this "zeroth-order" approximation, Q_{nm} becomes

$$Q^{(0)}_{nm} = A \left(\frac{2\pi}{\Delta k_\perp}\right)^2 \frac{\exp[i\,\boldsymbol{k}_\perp \cdot \boldsymbol{r}_{nm}]}{(P_{\rm los} + P_{{\rm N},n})\;(P_{\rm los} + P_{{\rm N},m})}, \qquad (26)$$

where $P_{\rm los} = P(k_\parallel, 0)$ and

$$A \;=\; \left(\sum_{n\,m} (P_{\rm los} + P_{{\rm N},n})^{-1}\;(P_{\rm los} + P_{{\rm N},m})^{-1}\right)^{-1}, \qquad (27)$$

$$\approx\; \left(\sum_{n} (P_{\rm los} + P_{{\rm N},n})^{-1}\right)^{-2}. \qquad (28)$$

The last line assumes that the noise is uncorrelated between sightlines and neglects the terms with $m = n$, whose fractional contribution vanish as $N \to \infty$. The final simplifications makes it apparent that this weighting is identical to that derived in Section 2.2, where we had implicitly made the same assumptions. The estimator in Section 2.2 is biased at the $N^{-1/2}$ level. Subtracting the term ${\rm tr}[\boldsymbol{Q}^{(0)}\boldsymbol{C}]$ from $\sum_{nm} Q_{nm}\,\tilde{\delta}_{{\rm F},n}\tilde{\delta}_{{\rm F},m}$ in equation (23) corrects for this bias.

Ignoring the ${\rm tr}[\boldsymbol{Q}^{(0)}\boldsymbol{C}]$ correction, the corresponding estimator for measuring the correlation function $\xi(\boldsymbol{r})$ follows from taking the Fourier transform of $Q^{(0)}_{nm}\,\tilde{\delta}_{{\rm F},n}\tilde{\delta}_{{\rm F},m}$ and, with the approximation in equation (28), is

$$\widehat{\xi^{(0)}}(r) = \sum_{nm}^{\text{fixed } \boldsymbol{r}} \left[w_n(x) \star \delta_{\rm F}(\boldsymbol{x}_n)\right] \star \left[w_m(x) \star \delta_{\rm F}(\boldsymbol{x}_m)\right], \qquad (29)$$

where $\star$ signifies a convolution along the line-of-sight. This estimator is identical to that derived in a different manner in Section 2.4.

Figure 4 plots the fractional difference between the variance of $\widehat{P}_F^{(0)}$ and that of the minimum variance quadratic estimator $\widehat{P}_F^{\rm QE}$, both calculated using equation (25) but with the appropriate $\boldsymbol{Q}$. The top panel fixes $k_\parallel$ at 0.1 Mpc^{-1} and varies $k_\perp$, and the bottom panel sets $k_\parallel = k_\perp$. The solid curves in the top panel show the fractional difference in the zero-noise limit ($P_{{\rm N},n} = 0$), which is the case where the fractional difference is largest. In order of increasing thickness, the solid curves represent $\bar{n} = 10^{-3}$, 3×10^{-3} and 10^{-2} Mpc^{-2}. At the smallest $k_\perp$ shown in the top panel, the estimator variance for $\widehat{P}_F^{(0)}$ is only 6% larger for the case with $\bar{n} = 10^{-3}$ Mpc^{-2} (thin solid curve). This difference becomes 10% and 20% for $\bar{n} = 3 \times 10^{-3}$ Mpc^{-2} and 10^{-2} Mpc^{-2}, respectively (thicker solid curves), and it increases with $k_\perp$. These estimates for the fractional increase in variance do not include sample variance, which would further decrease the plotted fraction (especially for the highest $\bar{n}$ and smallest k).

The red dashed curves in the top panel of Figure 4 are the same as the black solid curves except that they include noise. In particular, these curves assume a power-law luminosity function with slope -2, $z = 2.5$, $P_{{\rm N},n} = 1$ Mpc^{-1} where quasar n is at the limiting magnitude of the survey, and $S/N \propto$ flux. These choices result in $\bar{n}_{\rm eff}$ being reduced by a factor of 0.7 compared to the noiseless case, and the red dashed curves are suppressed relative to the black solid curves by a comparable factor. We find a similar correspondence for other noise models.

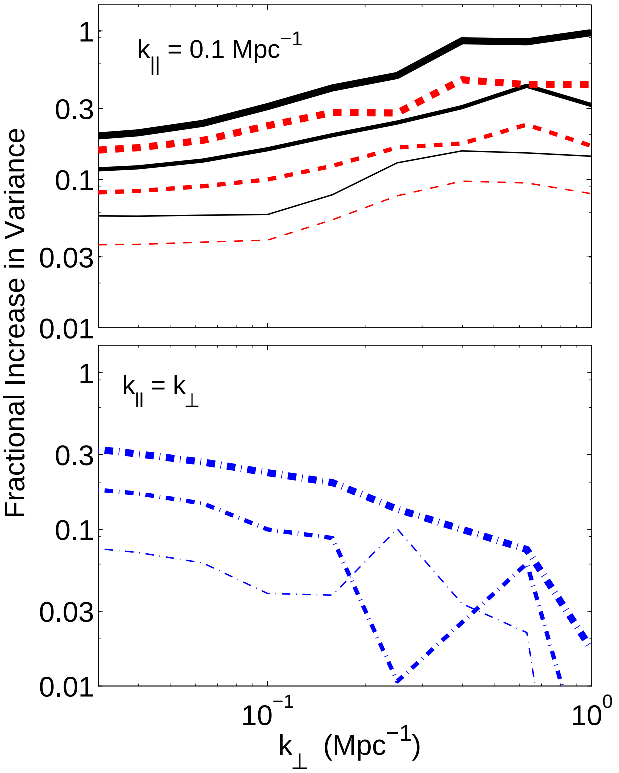

Figure 4. Top panel: Fractional increase in variance for the $\widehat{P}_{\rm F}^{(0)}$ estimator relative to the minimum variance quadratic estimator, $\widehat{P}_{\rm F}^{\rm QE}$, as a function of $k_\perp$. In order of increasing thickness, the curves are for $\bar{n} = 10^{-3}$, 3×10^{-3}, and 10^{-2} Mpc^{-2}. The black solid curves are the noiseless case, and the red dashed curves include noise as described in the text. These calculations assume $k_\parallel = 0.1$ Mpc^{-1}, are performed in a 200 Mpc × 200 Mpc region, and do not include sample variance. Bottom Panel: The blue dot-dashed curves are the same as the red dashed curves in the top panel but with $k_\parallel = k_\perp$.

One can improve upon our crude diagonal approximation for $\boldsymbol{C}$ in an iterative manner. In particular, maintaining only the diagonal elements in $\boldsymbol{C}$ when inverting can be thought of as the lowest order approximation for $\boldsymbol{C}^{-1}$ in equation (24). There are iterative methods that can be applied to achieve higher order corrections (such as the Jacobi method, Neumann iteration, or the Gauss-Seidel method). We advocate for the Gauss-Seidel method here because it is guaranteed to converge since $\boldsymbol{C}$ is Hermitian, and we have indeed verified that it converges. This method yields after the $i^{\rm th}$ iteration the estimate for the inverse of $\boldsymbol{C}$ given by

$$[\boldsymbol{C}^{(i)}_{\rm GS}]^{-1} = \boldsymbol{L}^{-1}\left(\boldsymbol{I} - \boldsymbol{U}\,[\boldsymbol{C}^{(i-1)}_{\rm GS}]^{-1}\right), \qquad (30)$$

where $\boldsymbol{L}$ and $\boldsymbol{U}$ are the lower and strictly upper diagonal

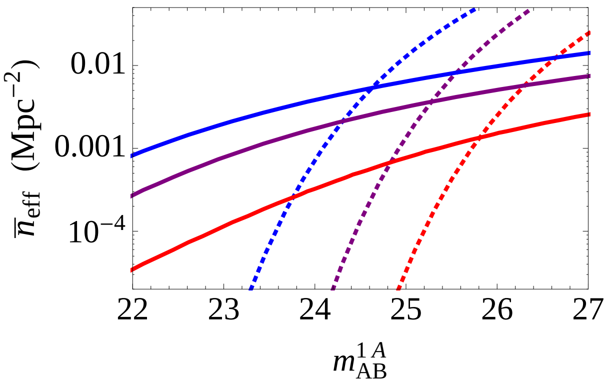

Figure 5. $\bar{n}_{\rm eff}$ as a function of $m_{\rm AB}^{1A}$ for a Lyα forest survey that uses either galaxies or quasars. These curves assume $m_{\rm AB}^{\rm lim} = m_{\rm AB}^{1A}$, where $m_{\rm AB}^{\rm lim}$ is the survey limiting magnitude. The solid curves show $\bar{n}_{\rm eff}$ for quasars at $z = 2, 3,$ and 4 (from top to bottom). The dashed curves show this quantity for galaxies at $z = 2, 3,$ and 4 (from left to right). Here, $\bar{n}_{\rm eff}$ is calculated at $k_\parallel = 0.1$ Mpc^{-1}, but its $k_\parallel$ dependence is extremely weak.

components of $\boldsymbol{C}$.[4] We initialize the iteration with $C^{(0)}_{{\rm GS},nm} = C_{nm}\,\delta^m_n$ so that the zeroth order iteration yields $Q^{(0)}$ (eqn. 26) and the associated estimator, $\widehat{P}^{(0)}_{\rm F}$.

The first order $\boldsymbol{Q}$ down-weights sightlines that fall near one another. It is instructive to write the first order result using the simpler Jacobi iteration method. In the Jacobi method, the first order approximation for the $\boldsymbol{C}^{-1}$ is

$$[\boldsymbol{C}^{(1)}_{\rm J}]^{-1} = [\boldsymbol{C}^{(0)}_{\rm J}]^{-1}(2\,\boldsymbol{C}^{(0)}_{\rm J} - \boldsymbol{C})\,[\boldsymbol{C}^{(0)}_{\rm J}]^{-1}. \tag{31}$$

Being cavalier with the normalization, the next order weights in the Jacobi iteration scheme are given by

$$Q^{(1)}_{{\rm J},nm} = \tilde{w}_n \tilde{w}_m\, e^{i\boldsymbol{k}_\perp\cdot\boldsymbol{r}_{nm}} \left(1-\tilde{S}_n\right)\left(1-\tilde{S}^*_m\right), \tag{32}$$

where

$$\tilde{S}_i(\boldsymbol{k}) = \sum_{\forall j, j\neq i} \frac{P(k_\parallel, \boldsymbol{r}_{ij})}{P_{\rm los}(k_\parallel) + P_{{\rm N},j}} \exp[-i\boldsymbol{k}_\perp\cdot\boldsymbol{r}_{ij}]. \tag{33}$$

The average of S_i in a survey equals $P_{\rm F}(\boldsymbol{k})/[\bar{n}_{\rm eff}^{-1} P_{\rm los}]$. If the number of quasars within $|\boldsymbol{k}_\perp\cdot\boldsymbol{r}_{ij}| \lesssim 1$ of quasar i is larger than for other quasars, S_i will also be larger, suppressing the weight given to this quasar. However, if all quasars have a similar number of quasars within $|\boldsymbol{k}_\perp\cdot\boldsymbol{r}_{ij}| \lesssim 1$, then the S_i would not vary strongly among the sightlines and this correction would be of minimal importance.

4 SURVEY CONSIDERATIONS

4.1 Galaxies

The apparent flatness of the faint end of the quasar luminosity function does not facilitate using observations of quasars with $L < L_*$ to obtain a denser sample of Lyα forest spectra. However, galaxies can be used to supplement $\bar{n}_{\rm eff}$ in very deep surveys. The dashed curves in Figure 5 represent $\bar{n}_{\rm eff}$ for a spectroscopic galaxy survey at $z = 2, 3,$ and 4 (from left to right), assuming that $m_{\rm AB}^{\rm lim} = m_{\rm AB}^{1A}$. These curves are calculated from the luminosity function of Lyman-break galaxies in Bouwens et al. (2007) for $z \geqslant 3$, assuming that there is no evolution in the ultraviolet luminosity function between $z = 2$ and 3 (Reddy & Steidel 2009). The galaxy magnitudes are for rest-frame 1600 Å. The solid curves are the same but for a quasar survey at these redshifts. Galaxies begin to aid the sensitivity when $m_{\rm AB}^{1A} \approx 24$ at $z = 2$, when $m_{\rm AB}^{1A} \approx 25$ at $z = 3$, and when $m_{\rm AB}^{1A} \approx 26$ at $z = 4$. Including galaxies results in a dramatic increase in $\bar{n}_{\rm eff}$ once a survey pushes below these stringent magnitude thresholds (Table 4).

The non-smooth continuum of galaxies' spectra will add an additional source of noise, but will not bias the estimate of $P_{\rm F}$ as long as correlations within a single sightline are neglected and the mean continuum can be removed (and even these mistakes can be isolated; Section **??**). If the power in galaxies' continuum spectra is smaller than $\bar{n}/\bar{n}_{\rm eff} \times P_{\rm los}$, using galaxies in addition to quasars has the potential to significantly improve a survey's sensitivity.

4.2 Cross Correlation

It is also possible to cross-correlate Lyα forest skewers with some other tracer of large-scale structure. This could be done using a galaxy survey or with the quasars in the Lyα survey themselves. We express the overdensity field of these tracers as δ_g and their average 3D number density as $\bar{n}_{\rm g,3D}$. The average cross power between these signals is $\widehat{P}_{\rm Fg}(k) \equiv \langle\tilde{\delta}(k)\,\tilde{\delta}_{\rm g}(k)\rangle$ and the variance on an estimate of this signal is

$$\text{var}[\widehat{P}_{\rm Fg}(\boldsymbol{k})] = P_{\rm Fg}(\boldsymbol{k})^2 + P_{\rm tot}(\boldsymbol{k})\,\left(P_{\rm g}(\boldsymbol{k}) + \bar{n}^{-1}_{\rm g,3D}\right), \tag{34}$$

where $P_{\rm Fg}$ is the cross-power spectrum between $\delta_{\rm F}$ and $\delta_{\rm g}$ and $P_{\rm g}$ is the auto-power of $\delta_{\rm g}$. If we again assume the weight factorizes, weights given by equation (11) also maximize the S/N in this case.

There are three scenarios for which the cross power could be an interesting measurement: (1) the cross-correlation is more sensitive to $P_\delta^{\rm lin}$ than the Lyα forest auto-correlation, (2) it could be used to separate different contributions to the Lyα flux power or to measure the bias of the galaxies, or (3) it could be a systematic check for both the Lyα and the other survey.

Cross correlation is more sensitive than the Lyα forest auto-correlation when

$$n_{\rm g,3D} > \bar{n}_{\rm eff}\, P_{\rm los}^{-1}(k_\parallel)\, b^2/b_g^2, \tag{35}$$

where we have assumed that the shot noise and aliasing terms are the dominant sources of variance for both surveys. For $\bar{n}_{\rm eff} = 10^{-3}$ Mpc^{-2}, $b = 0.3$, and $P_{\rm los} = 0.5$ (where the latter two numbers are characteristic of the $z = 3$ Lyα forest), condition (35) becomes $n_{\rm g,3D} > 2 \times 10^{-5}$ Mpc^{-3} for a tracer with $b_g = 3$. In addition, at scales where the other survey is instead sample variance-limited, the S/N in cross correlation will be at least as large as in the Lyα forest survey's auto-correlation.

Inequality (35) requires a number density that is a factor of ≈ 10 higher than the 3D quasar density that BOSS aims to measure. However, even though the S/N in the cross correlation is then a factor of ≈ 10 below the S/N in the Lyα forest auto correlation, this should be still be sufficient to measure the quasar bias (and to much better precision than is possible with the quasar auto-power). A more promising route may be to cross correlate with a separate high-redshift galaxy survey. As an example, the Hobby-Eberly Telescope Dark Energy Experiment (HETDEX)[5] aims to find spectroscopically 1 million Lyα emitters between $1.8 < z < 3.8$ over 200 deg^2, which works out to an average density of $\bar{n}_{\rm g,3D} \approx 2 \times 10^{-4}$ Mpc^{-3} and easily satisfies criterion (35) when correlated with a BOSS-like Lyα forest survey.

---

[4] The inversion of the lower diagonal matrix $\boldsymbol{L}$ requires at most N^2 operations.

If criterion (35) is satisfied, the S/N is greater than or equal for the galaxy survey's own auto-power compared to its cross power with the Lyα forest survey. Thus, the cross power may be most interesting as a systematics check as well as to investigate the sources of fluctuations that contribute to $\delta_{\rm F}$.

4.3 Survey Strategy

Figure 6 illustrates the survey strategy tradeoffs in the sky coverage versus magnitude-limit plane for $z = 2.5$. The black solid curves in both panels are contours of constant sensitivity to $P_{\rm F}$ for a quasar survey. The black long-dashed curves are for a survey with quasars plus galaxies and at the same sensitivity as the black solid curves. The framed labels associated with each solid curve show the S/N of a measurement of $P_{{\rm F},\mu}$ in a k-space shell of width $\Delta k/k = 0.2$. This shell is centered on $k = 0.1$ Mpc^{-1} (left panel) or on $k = 0.5$ Mpc^{-1} (right panel).

The red short-dashed curves are contours of constant $\mathcal{M} \equiv$ Survey Area/(limiting flux)^2. For fixed duration of the observing campaign on an instrument, $\mathcal{M}$ is a constant that does not depend on the survey strategy (assuming that the limiting flux decreases as exposure time on a field to the 1/2 power). The curves in Figure 6 correspond to $\mathcal{M}$ of 10^{-2}, 10^{-1}, 1, 10, and 100 times $\mathcal{M}$ for the BOSS survey, whose approximate specifications are represented by the filled circle ($m_{\rm AB}^{\rm lim} = m_{\rm AB}^{1A} = 22$, $\mathcal{A} = 8000$ deg^2). For the present specifications of BigBOSS, $\mathcal{M}$ is ≈ 7 times larger than for BOSS. McDonald et al. (2005) suggested an extension of the WFMOS spectroscopic galaxy survey to include quasars (Glazebrook et al. 2005) and would reach $m_{\rm AB}^{\rm lim} \approx 24.5$. WFMOS is a proposed 300 deg^2 survey on the 8 m Subaru telescope. This hypothetical quasar survey would be only moderately more sensitive than BOSS at the considered scales.

The left panel in Figure 6 argues that the survey strategy for BOSS is close to optimal in the sense of minimizing var$[\widehat{P}_{\rm F}]$ at $k = 0.1$ Mpc^{-1}: The sensitivity contours have a similar shape to the red short-dashed curve that intersects the BOSS point. Similar conclusions hold for the sensitivity at $k = 0.5$ Mpc^{-1} (right panel, Fig. 6). In general, the optimal strategy for a quasar survey corresponds to integrating deep enough on each field (and having enough fibers/slits) to reach quasars with luminosities of $\approx L_*$, where L_* is the characteristic luminosity in the quasar luminosity function. While Figure 6 assumes that $m_{\rm AB}^{\rm lim} = m_{\rm AB}^{1A}$, the vertical axis can also be thought of as a function of $\bar{n}_{\rm eff}$, noting that $m_{\rm AB}^{1A} = m_{\rm AB}^{\rm lim} = \{22,\ 24,\ 26\}$ corresponds to $\bar{n}_{\rm eff} = \{0.5,\ 3,\ 7\}$ Mpc^{-2} at $z = 2.5$ for a survey that only includes quasars (Table 4).

At the considered wavevectors and fixed $\mathcal{M}$, the sensitivity to $P_{\rm F}$ depends little on a survey's depth, at least for $m_{\rm AB}^{\rm lim} < 23$. This result owes to the shape of the bright end of the luminosity function. However, we favor the survey strategy that is deep enough to reach down to L_* because deeper quasar surveys will be able to better handle systematics (because they have higher S/N per mode) and they rely less on the noise decreasing as the square root of the number of modes (which can be invalidated by non-Gaussian effects; Meiksin & White 1999). Technically, a shallower but wider survey strategy will be more sensitive at smaller k than shown here because it samples more modes.

The sensitivity gains of a deep survey that uses the spectra from galaxies in addition to quasars are not always significant. Compare the long-dashed and solid curves in Figure 6. The denser sample of sightlines that galaxies provide is more helpful at $k = 0.5$ Mpc^{-1} compared to $k = 0.1$ Mpc^{-1} (Fig. 6), but the sensitivity is never significantly improved over a quasar survey covering a larger fraction of the sky (fixing $\mathcal{M}$). Thus, the small gains in sensitivity from including galaxies may not be sufficient to outweigh the added difficulty of removing the galaxies' more complex continua.

It is much more difficult to measure $P_{\rm F}$ at $z = 4$. However, some of the science drivers of a measurement at $z = 4$ may not require a precision measurement, and even an order unity measurement at this redshift would be interesting (McQuinn et al. 2010). Figure 7 quantifies the prospects for a measurement of $P_{\rm F}$ at $z = 4$. A spectroscopic quasar survey reaching $m_{\rm AB}^{1A} = 22$ (the same depth as BOSS) would be able to achieve $S/N \approx 3$ at $k = 0.1$ Mpc^{-1} in shells of width $\Delta k/k = 0.2$. A survey reaching $m_{\rm AB}^{1A} = 23$ would achieve $S/N \approx 10$.

Thus far we have not discussed how to select quasars as a function of redshift: A survey with a multi-object spectrograph would in practice have to decide whether to obtain the spectra of a quasar at $z = 2$ or $z = 3$, and this choice would affect its sensitivity. In the aliasing noise-limited regime, the total $[S/N]^2$ on a scale is $\propto \int dz\, \bar{n}_{\rm eff}^2$. Because this integral is quadratic in $\bar{n}_{\rm eff}$, maximizing $\bar{n}_{\rm eff}$ over a limited interval in redshift also maximizes the S/N. However, once sample variance limits a survey at the scales of interest, it makes sense to broaden the redshift width of the quasar selection function. In practice, a small fraction of slits or fibers per degree compared to the total number used in modern spectroscopic galaxy surveys is required to select all $z > 2$ quasars down to a survey's limiting flux ($\approx 10 - 100$). Therefore, the decision of which quasars to target may often be moot.

4.4 Cosmological and Astrophysical Constraints

Ultimately, one wants to use 3D Lyα forest measurements to constrain cosmological parameters, ionizing background models, and the reionization history. This section briefly discusses how well these quantities can be constrained. To do so, we calculate the Fisher matrix defined as (e.g., Tegmark et al. 1997)

$$F_{ij} \equiv - \left\langle \frac{\partial^2 \ln \mathcal{L}}{\partial \lambda_i \partial \lambda_j} \right\rangle = \frac{1}{2} \sum_{\forall \boldsymbol{k}} P_{\rm tot}(\boldsymbol{k})^{-2} \, \frac{\partial P_{\rm F}(\boldsymbol{k})}{\partial \lambda_i} \frac{\partial P_{\rm F}(\boldsymbol{k})}{\partial \lambda_j}, \tag{36}$$

where $\mathcal{L}$ is the likelihood of the model given the data, the second equality assumes Gaussianity, and the λ_i are the pa-

---

[5] `http://www.as.utexas.edu/hetdex/`

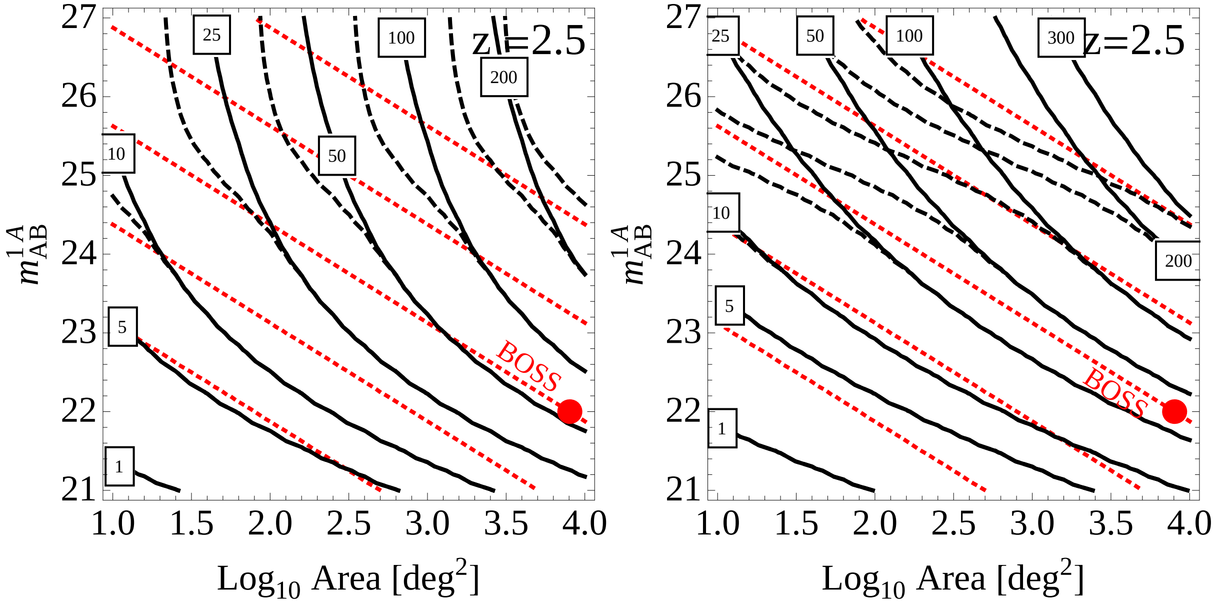

Figure 6. Sensitivity to $P_{\rm F}$ at $k = 0.1$ Mpc^{-1} (left panel) and $k = 0.5$ Mpc^{-1} (right panel) as a function of survey area and limiting magnitude, where we have assumed that $m_{\rm AB}^{\rm lim} = m_{\rm AB}^{1A}$, $L \approx 500$ Mpc, and $z = 2.5$. The black solid curves are contours of constant sensitivity for a quasar survey and the partially overlapping black long-dashed curves are at the same sensitivity but for a survey that also includes galaxies down to the same magnitude. The framed labels associated with each solid curve quote the S/N on $P_{{\rm F},\mu}$ in bins of width $\Delta k/k = 0.2$. The red circle represents the approximate specifications of the BOSS quasar survey. The red short-dashed curves are contours of constant $\mathcal{M} \equiv$ Survey Area/(limiting flux)^2. For fixed duration of the observing campaign on an instrument, $\mathcal{M}$ does not depend on the survey strategy. The red short-dashed curves correspond to 10^{-2}, 10^{-1}, 1, 10, and 100 times the $\mathcal{M}$ of the BOSS survey.

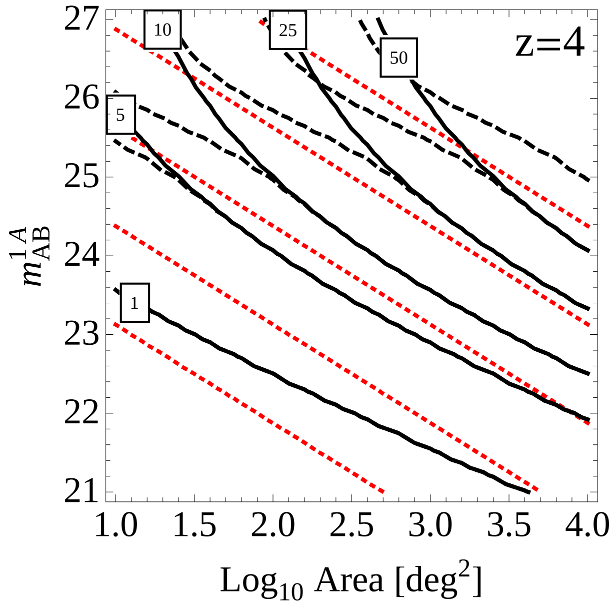

Figure 7. Same as Figure 6 but at $z = 4$ and for $k = 0.1$ Mpc^{-1}

rameters we want to constrain. The summation is over all independent modes. Given this parameter set, the forecasted $1\,\sigma$ uncertainty on the parameter λ_i is $[\boldsymbol{F}^{-1}]_{ii}^{1/2}$.

We perform our calculations for a survey at $z = 2.5$ and with a base parameter set that is given by the amplitude of $P_{\rm F}$, its tilt and running with pivot at $k = 0.1$ Mpc^{-1}, the redshift-space distortion parameter g, the angular diameter distance D_A, and the Hubble expansion parameter $H(z)$. The former four parameters we treat as nuisance parameters. Finally, we omit the first three line-of-sight modes since these are likely to be contaminated by continuum and mean flux errors, as motivated in McDonald et al. (2005). The results depend weakly on this assumption.

Some of the constraint from our base parameter set on $D_A(z)$ and $H(z)$ owes not just to the baryon acoustic oscillation (BAO) features, but also to Alcock-Paczynski type effects. Since it might be the case that the continuum under the BAO is contaminated by other effects, one may not want to use information that derives from this broad-band power. To guarantee that the constraint owes to the BAO features, we also provide a more conservative estimate in which we subtract from $P_{\rm F}$ a flux power spectrum that does not include the BAO prior to calculating $\partial P_{\rm F}/\partial\lambda_i$. In what follows, we quote both constraints.

Figure 8 plots the constraints on $D_A(z)$ and $H(z)$ at $z = 2.5$ as a function of $\bar{n}_{\rm eff}$. The curves are for a survey with $L = 1$ Gpc and $\mathcal{A} = 10^4$ deg^2. The thin curves represent the optimistic predictions using the fiducial estimate for $\boldsymbol{F}$, and the thick curves are from the more conservative estimate that subtracts the continuum underneath the BAO. A BOSS-like survey with $A \approx 10^4$ deg^2, $L \approx 1$ Gpc, and $\bar{n}_{\rm eff} \approx 5 \times 10^{-4}$ Mpc^{-2} ($m_{\rm AB}^{\rm lim} = 22$) can constrain D_A at $z = 2.5$ to fractional precision of $\approx 0.02 - 0.03$, and one like BigBOSS where $\bar{n}_{\rm eff} \approx 1.5 \times 10^{-3}$ Mpc^{-2} ($m_{\rm AB}^{\rm lim} = 23$) to a fractional precision of $\approx 0.010 - 0.015$.[6] The constraints we found on $H(z)$ and $D_A(z)$ are comparable to those in

---

[6] The angular diameter distance is generally better constrained in surveys of large-scale structure for geometric reasons. However, the strength of redshift-space distortions in Lyman-α forest sur-

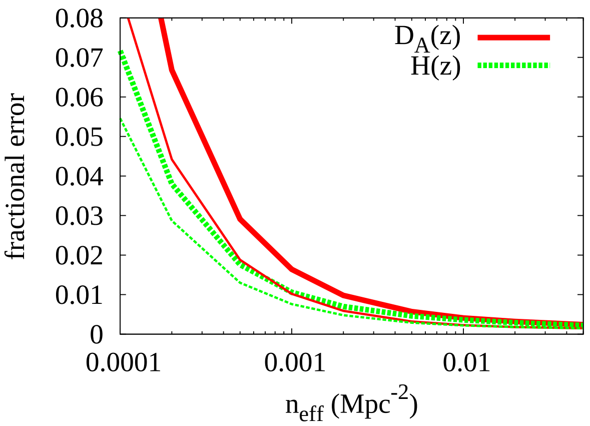

Figure 8. Fractional error on the angular diameter distance and Hubble expansion parameter at $z = 2.5$ as a function of $\bar{n}_{\rm eff}$. The curves are calculated for a survey with $L = 1$ Gpc and $\mathcal{A} = 10^4$ deg^2, but their amplitude scales as $(L\,\mathcal{A})^{-1/2}$. The thick curves represent a conservative estimate, and the thin curves represent an optimistic estimate, as described in the text.

McDonald et al. (2005). Figure 8 allows one to estimate the sensitivity of a survey of arbitrary volume since the error scales as $(L\,\mathcal{A})^{-1/2}$.

The precision at which $D_A(z)$ is measured from the Lyα forest provides a constraint on cosmological parameters, such as the curvature space density, Ω_k. In particular, a measurement of $D_A(z)$ at $z \approx 2.5$ differenced with its value at $z = 1100$ – which is tightly constrained by cosmic microwave background measurements – provides a robust measure of Ω_k (at least under the assumption that dark energy is negligible at $z > 2$). If the principle uncertainty is in the angular diameter distance to the Lyα forest, as is likely to be the case, this translates to the constraint $\delta\Omega_k \approx 0.26\,\delta D_A/D_A$ at $z = 2.5$. Thus, a measurement of Ω_K to the precision $\approx 3\times10^{-3}$ is possible with a BigBOSS-like survey, an improvement over the present-day limit of $|\Omega_K| < 1\%$ that assumes ΛCDM (Komatsu et al. 2011).

Ionizing background and temperature fluctuations can also be measured in a 3D Lyα forest analysis. To gauge how well these contributions can be constrained, in addition to our base set of parameters, we also include a scale-invariant contribution to $P_{\rm F}$ that scales as k^{-3}, a component that scales as k^{-2}, and a white component that scales as k^0 (all spherically symmetric). The contribution that scales as k^{-2} is motivated by the expected scaling for intensity fluctuations at wavevectors larger than one over the mean free path of 1 Ry photons (essentially all accessible k at $z = 2.5$; McQuinn et al. 2010), and the scale-invariant form approximates some of the temperature models in McQuinn et al. (2009). We have set the normalization of each component to be such that it is the same as in our simple model for $P_{\rm F}$ at $(k_\parallel, k_\perp) = (0.1, 0)$ Mpc^{-1}. This normalization is comparable to that of physically motivated models for temperature and intensity fluctuations presented in McQuinn et al. (2010). For a BOSS-like survey, the normalization of the k^{-2} component is constrained to 10%, the normalization of the k^{-3} component to 3%, and the normalization of the white component to 0.03%. We have also examined the constraints for a shallower survey with $\bar{n}_{\rm eff} = 10^{-4}$ Mpc^{-2}, $L = 1$ Gpc, and $\mathcal{A} = 10^4$ deg^2 (as could be obtained by BigBOSS if it targets $z \approx 4$). This case yields 40% and 12% constraints on the k^{-2} and k^{-3} components, respectively.

5 SYSTEMATIC CONCERNS

This section briefly discusses systematics that are known to be serious for line-of-sight Lyα forest analyses in the context of 3D analyses. We focus on continuum fitting errors, mistakes in the mean flux estimate, and the damping wings of high column-density absorption systems. A pervading theme of this discussion is the advantages of a 3D analysis in dealing with these systematics.

Measurements of $P_{\rm F}(\boldsymbol{k})$ must remove the structure in the quasar continuum in order to achieve an unbiased measurement of $\delta_{\rm F}$. However, continuum fluctuations do not have to be removed perfectly since they only act as a source of noise as long as the error in subtracting them does not correlate between different sightlines (Viel et al. 2002). (In addition, any systematic error introduced in this removal can be excised at a later step as described below.) The extra continuum power that does not correlated between the different continua, $P_{\rm los}^{\rm cont}$, would contribute an extra term in $P_{\rm tot}$ equal to $\bar{n}^{-1}P_{\rm los}^{\rm cont}$, where again $\mathrm{var}[\widehat{P}_{\rm F}] = 2\,P_{\rm tot}^2$, assuming that line-of-sight correlations are omitted. At $k_\parallel > 0.1$ Mpc^{-1}, McDonald et al. (2005) found that the mean quasar continuum had power equal to $< 1\%$ of $P_{\rm F}$. Thus, continuum fluctuations will not significantly increase the variance (by increasing the amount of line-of-sight power) at $k_\parallel > 0.1$ Mpc^{-1}, but could be of more importance at smaller wavevectors.

The redshift evolution of the mean flux $\langle F\rangle$ – which can be quite substantial over $L \sim 500$ Mpc – can also contribute spurious power at small $k_\parallel$. Accurately estimating $\langle F\rangle$ is known to be crucial for interpreting line-of-sight measurements. It also has the potential to be an even more important issue for 3D measurements, where the size of fluctuations one aims to measure are significantly smaller. For example, the dimensionless 3D power at $k = 0.1$ Mpc^{-1} is $\Delta_{\rm F}^2 \sim 10^{-4}$, which would require that the mean flux evolution be known to precision $\ll 10^{-2}$ in order to not bias the measurement of $P_{\rm F}$ averaged over a shell centered at this k. (Although, the power from the mean flux error would primarily impact purely line-of-sight modes.) Traditionally, $\langle F\rangle$ is estimated by averaging the transmission at redshift z from all the sightlines in a survey. Given N sightlines and if we want to estimate $\langle F\rangle$ in a region of size $\Delta\chi$, then this quantity can be estimated to precision

$$\sigma_{\langle F\rangle}^2 \approx \frac{(P_{\rm los} + P_{\rm N})}{\Delta\chi\, N}. \tag{37}$$

Noting that the numerator on the right hand side is $\mathcal{O}(1)$, it takes 1000 Lyα forest spectra to constrain $\sigma_{\langle F\rangle}$ in a 10 Mpc region to reach the threshold of 10^{-2}, which is much less than the total number of Lyα forest spectra in BOSS or BigBOSS.

However, stacking to obtain $\langle F\rangle$ requires that no mis-

---

veys enhances the line-of-sight power, boosting the sensitivity to $H(z)$.

takes are made. For example, a systematic error in the continuum subtraction can lead to a biased estimate of $\langle F \rangle$. Fortunately, in 3D analyses, one has the ability to discern these effects at a later step in the analysis. A mean flux error will principally lead to spurious power that only affects purely line-of-sight modes. Thus, one can simply throw away the first X modes along the origin in the analysis to remove these errors. In addition, any bleeding to other modes owing to the complicated survey window may not be so worrisome because the error in the mean flux is at the percent-level such that the total variance contributed by this error is comparable to Δ_F^2. For example, the power in this bleeding for the $P_{\rm F}$ estimator described in Section 2 is suppressed by $N^{-1/2}$ relative to the mean flux-error power in the line-of-sight modes. Errors in the mean flux also enter in convolution with $P_{\rm F}$, but this effect is less of a contaminant again given that the sizes of these errors are percent-level over $x \sim 100$ Mpc. For such an error, a percent of $\delta_{\rm F}$ is smoothed over $\Delta k_{\parallel} \sim x^{-1}$.

Another systematic is the damping wing absorption from Lyman-limit and damped Lyman-α systems. This absorption is generally not included in simulations of the Lyα forest, and it arises from dense, self-shielding systems that are not captured properly in almost all cosmological simulations anyway (Katz et al. 1996; McQuinn et al. 2011). These systems could alter the line-of-sight power at the 10% level at $k_{\parallel} \approx 0.1$ Mpc^{-1} (McDonald et al. 2005), and at the 100% level at $k_{\parallel} = 0.01$ Mpc^{-1} (Appendix B).

Appendix B discusses a simple model for the power in these systems that qualitatively reproduces the numerical estimates for the effect of damping wings in McDonald et al. (2005) and that allows us to estimate their impact for 3D analyses. We show that much of the line-of-sight power from damping wings arises from the uncorrelated (shot) component of their power. The shot contribution to the 3D flux power from damping wings is much smaller than in the line-of-sight flux power spectrum, and, thus, the role of damping wings as a contaminant is reduced (Appendix B). However, we show that they still are likely to add power to $P_{\rm F}$ at the tens of percent level.

6 CONCLUSIONS

This paper studied issues relevant to upcoming 3D Lyα forest surveys. We derived a simple formula for how to optimally weight sightlines with varying S/N levels. We found that simply weighting sightline i by $(1+\sigma_{{\rm N},i}^2/\sigma_{\rm los}^2)^{-1}$, where $\sigma_{{\rm N},i}^2$ and $\sigma_{\rm los}^2$ are respectively the variance of the noise and of the Lyα forest normalized flux in a 10 Mpc region or greater, performs nearly as well as the weights specified by the minimum variance quadratic estimator. These weights should be simple to apply to data even in the presence of real-world complications. We derived an expansion (which involves only matrix multiplications) that converges to the minimum variance quadratic estimator and for which our suggested weights are the lowest order contribution. We showed that the next term in this expansion has the intuitive behavior of suppressing the contribution from quasars that have an overabundance within $r_\perp \lesssim k_\perp^{-1}$.

We showed that the sensitivity of a spectroscopic survey to the Lyα forest flux power spectrum can be quantified by just a single number, $\bar{n}_{\rm eff}$ – a noise-weighted number density of sources on the sky – so that ${\rm var}[\widehat{P}_{\rm F}(\boldsymbol{k})] = 2\,(P_{\rm F}+\bar{n}_{\rm eff}^{-1}\,P_{\rm los})^2$. While this number technically depends on the $k_\parallel$ of the wavevector in question, in practice this dependence is extremely weak at $k_\parallel < 0.5$ Mpc^{-1} because $P_{\rm los}$ is almost constant over these wavevectors. These are the same wavevectors at which 3D Lyα surveys have the potential to derive competitive constraints on cosmological parameters (McDonald & Eisenstein 2007) and that are the most interesting for studying astrophysical sources of fluctuations in the Lyα forest (McQuinn et al. 2010).

We calculated $\bar{n}_{\rm eff}$ as a function of survey specifications for both quasar and galaxy surveys at different redshifts. For quasar surveys, it is difficult to achieve significantly higher $\bar{n}_{\rm eff}$ than $\approx 3 \times 10^{-3}$ Mpc^{-2} (or ≈ 30 per deg^2) at any redshift owing to the shallowness of the faint end of the quasar luminosity function. A survey with $\bar{n}_{\rm eff} = 3 \times 10^{-3}$ Mpc^{-2} is aliasing noise-limited at $k > 0.1$ Mpc^{-1}. In this limit, the S/N on $P_{\rm F}$ scales linearly with $\bar{n}_{\rm eff}$.

This paper also discussed survey strategy tradeoffs. The previous results allowed for a simpler characterization of these tradeoffs than in prior studies. We showed that a survey's sensitivity to the flux correlation function is maximized with the strategy of integrating on each field just long enough to achieve $S/N \approx 2$ in a 1 Å pixel for an L_* quasar. However, we found that a shallower strategy but covering a wider field formally obtains a comparable sensitivity to $P_{\rm F}$ at $0.1 < k < 0.5$ Mpc^{-1} (but with lower S/N per mode), and it could even be more sensitive at smaller k. While such low S/N values may cause worry regarding the efficacy of continuum subtraction, we further argued that continuum removal errors as well as other systematics are likely to be less problematic in 3D correlation analyses compared to in 1D. Lastly, we found that using the spectra from galaxies never results in the sensitivity being vastly improved over a shallower but wider survey that only incorporates quasars.

We quantified as a function of $\bar{n}_{\rm eff}$ a survey's sensitivity to the angular diameter distance at the mean redshift of the survey, to the Hubble expansion rate, and to the sources of fluctuations that do not trace density. Upcoming surveys have the potential to measure $D_A(z)$ and $H(z)$ to $1-2\%$, which would substantially improve constraints on the curvature of space-time and early dark energy models. These surveys also have the potential to detect other contributions to the flux power at the percent-level and, thereby, constrain the level of temperature fluctuations relic from reionization processes and the properties of the extragalactic ionizing background.

We thank Anže Slosar for helpful comments on this manuscript, and Shirley Ho, Adam Lidz, Nic Ross, Jennifer Yeh, and Matias Zaldarriaga for useful discussions. MM is supported by the NASA Einstein Fellowship.

REFERENCES

Becker G. D., Bolton J. S., Haehnelt M. G., Sargent W. L. W., 2011, MNRAS, 410, 1096
Becker R. H., et al., 2001, AJ, 122, 2850
Bouwens R. J., Illingworth G. D., Franx M., Ford H., 2007, ApJ, 670, 928
Cooray A., Sheth R., 2002, physrep, 372, 1

Croft R. A. C., Weinberg D. H., Katz N., Hernquist L., 1998, ApJ, 495, 44
Fan X., et al., 2006, AJ, 132, 117
Feldman H. A., Kaiser N., Peacock J. A., 1994, ApJ, 426, 23
Glazebrook K., Eisenstein D., Dey A., Nichol B., The WF-MOS Feasibility Study Dark Energy Team 2005, astro-ph/0507457
Hamilton A. J. S., Rimes C. D., Scoccimarro R., 2006, MNRAS, 371, 1188
Hennawi J. F., Prochaska J. X., 2007, ApJ, 655, 735
Hopkins P. F., Hernquist L., Cox T. J., Di Matteo T., Robertson B., Springel V., 2006, ApJS, 163, 1
Katz N., Weinberg D. H., Hernquist L., Miralda-Escude J., 1996, ApJL, 457, L57+
Komatsu E., et al., 2011, ApJS, 192, 18
Lidz A., Faucher-Giguère C., Dall'Aglio A., McQuinn M., Fechner C., Zaldarriaga M., Hernquist L., Dutta S., 2010, ApJ, 718, 199
Liske J., Webb J. K., Williger G. M., Fernández-Soto A., Carswell R. F., 2000, MNRAS, 311, 657
McDonald P., Eisenstein D. J., 2007, PRD, 76, 063009
McDonald P., et al., 2005, ApJ, 635, 761
McDonald P., Miralda-Escudé J., Rauch M., Sargent W. L. W., Barlow T. A., Cen R., 2001, ApJ, 562, 52
McDonald P., Miralda-Escudé J., Rauch M., Sargent W. L. W., Barlow T. A., Cen R., Ostriker J. P., 2000, ApJ, 543, 1
McDonald P., Seljak U., Cen R., Bode P., Ostriker J. P., 2005, MNRAS, 360, 1471
McQuinn M., Hernquist L., Lidz A., Zaldarriaga M., 2010, arXiv:1010.5250
McQuinn M., Lidz A., Zaldarriaga M., Hernquist L., Hopkins P. F., Dutta S., Faucher-Giguère C., 2009, ApJ, 694, 842
McQuinn M., Oh S. P., Faucher-Giguere C., 2011, arxiv:1101.1964
Meiksin A., White M., 1999, MNRAS, 308, 1179
Meiksin A. A., 2009, Reviews of Modern Physics, 81, 1405
O'Meara J. M., Prochaska J. X., Burles S., Prochter G., Bernstein R. A., Burgess K. M., 2007, ApJ, 656, 666
Prochaska J. X., O'Meara J. M., Worseck G., 2010, ApJ, 718, 392
Reddy N. A., Steidel C. C., 2009, ApJ, 692, 778
Rimes C. D., Hamilton A. J. S., 2006, MNRAS, 371, 1205
Schaye J., Theuns T., Rauch M., Efstathiou G., Sargent W. L. W., 2000, MNRAS, 318, 817
Seljak U., et al., 2005, PRD, 71, 103515
Slosar A., Ho S., White M., Louis T., 2009, Journal of Cosmology and Astro-Particle Physics, 10, 19
Tegmark M., Taylor A. N., Heavens A. F., 1997, ApJ, 480, 22
Viel M., Lesgourgues J., Haehnelt M. G., Matarrese S., Riotto A., 2005, PRD, 71, 063534
Viel M., Matarrese S., Mo H. J., Haehnelt M. G., Theuns T., 2002, MNRAS, 329, 848
Weinberg D. H., Miralda-Escude J., Hernquist L., Katz N., 1997, ApJ, 490, 564
White M., 2003, in The Davis Meeting On Cosmic Inflation The Ly-a forest
White M., Pope A., Carlson J., Heitmann K., Habib S., Fasel P., Daniel D., Lukic Z., 2010, ApJ, 713, 383
Williger G. M., Smette A., Hazard C., Baldwin J. A., McMahon R. G., 2000, ApJ, 532, 77

APPENDIX A: POWER-LAW LUMINOSITY FUNCTIONS

We gain some insight into $\bar{n}_{\rm eff}$ by imagining that the distribution of observed fluxes for quasars in the target redshift range is a power law,

$$N(f)=\frac{N_0}{f_0}\left(\frac{f}{f_0}\right)^{-\alpha}\qquad(\alpha>1),\tag{A1}$$

and that we observe all quasars above a given flux limit $f_{\rm min}$:

$$N=N_0\int_{f_{\rm min}}^{\infty}\frac{df}{f_0}\left(\frac{f}{f_0}\right)^{-\alpha}=\frac{N_0}{\alpha-1}\left(\frac{f_{\rm min}}{f_0}\right)^{1-\alpha}.\tag{A2}$$

If observing conditions are such that the noise in the forest for a quasar at flux f_0 has $P_{N,n}=\sigma_0^2\,P_{\rm los}$ and scales as f^{-2} then

$$\bar{n}_{\rm eff}=n_0\int_{f_{\rm min}}^{\infty}\frac{(f/f_0)^{-\alpha}}{1+\sigma_0^2(f_0/f)^2}\,\frac{df}{f_0},\tag{A3}$$

where $n_0\equiv N_0/\mathcal{A}$. The behavior of the integral is controlled by $\Phi\equiv\sigma_0 f_0/f_{\rm min}$; for general α it can be written in terms of a hypergeometric function, however for integer α it reduces to elementary functions. For example, when $\alpha=2$

$$\bar{n}_{\rm eff}=n_0\,\frac{\tan^{-1}\Phi}{\sigma_0},\tag{A4}$$

while for $\alpha=3$

$$\bar{n}_{\rm eff}=n_0\,\frac{\ln(1+\Phi^2)}{2\sigma_0^2},\tag{A5}$$

and larger α give more negative powers of σ_0 and combinations of logs or arctangents.

A natural choice for f_0 is to make $\sigma_0=1$, so that n_0 is the number of quasars with $P_N=P_{\rm los}$ and Φ measures the minimum flux in units of the characteristic flux. It is easy to see that decreasing $f_{\rm min}$, i.e. increasing Φ, leads to larger $\bar{n}_{\rm eff}$ but that the gains are small once $\Phi\gg1$. For $f_{\rm min}\sim f_0$ $\bar{n}_{\rm eff}$ is n_0 up to a numerical constant of order unity, which reinforces the discussion in the text.

APPENDIX B: DAMPED ABSORBERS

As briefly discussed in Section 5, an important systematic for Lyα forest analyses is the damping wing absorption from Lyman-limit and damped Lyman-α systems. A simple halo-like model for the contribution to the 3D power of the absorption from these systems (applicable at scales larger than those affected by thermal broadening) is

$$\begin{aligned}P_{\rm F}^{\rm DLA}&=\widetilde{W}_2(k_\parallel)^2\,(b_{\rm DLA}+\mu^2)^2\,P_\delta^{\rm lin}(k),\\ P_{\rm los}^{\rm DLA}&=\widetilde{W}_1(k_\parallel)+\widetilde{W}_2(k_\parallel)^2\int\frac{dk_\perp}{2\pi}k_\perp(b_{\rm DLA}+\mu^2)^2\,P_\delta^{\rm lin}(k),\end{aligned}\tag{B1}$$

where we have neglected the shot noise component in $P_{\rm F}^{\rm DLA}$ because it is small on relevant scales (having amplitude ∼

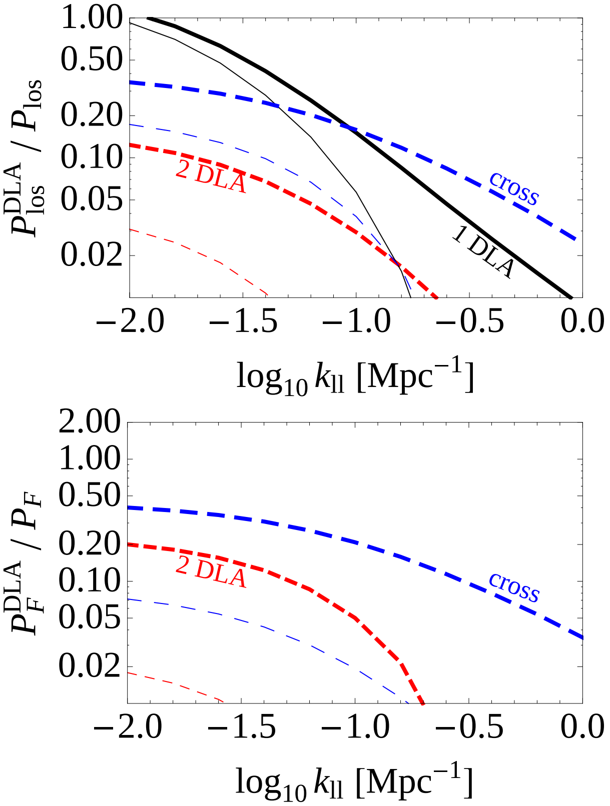

Figure B1. Ratio of different components of the power in damping wings to the Lyα forest power without damping wing absorption as a function of $k_\parallel$ and at $z = 2.5$. The top panel is for the line-of-sight power spectrum and the bottom is for the 3D power spectrum. The thick solid, thick short-dashed, and thick long-dashed curves are respectively the 1-DLA term, the 2-DLA term, and the Lyα forest–DLA cross power. The corresponding thin curves are the same but are calculated only using systems with $N_{\rm HI} > 10^{20}$ cm^{-2}.

1 Mpc^{-3}, which is much less than $b_{\rm DLA}^2 \, P_\delta^{\rm lin}$) and

$$\widetilde{W}_1(k_\parallel) = \int dN_{\rm HI} \, \frac{\partial^2 \mathcal{N}}{\partial\chi \partial N_{\rm HI}} \tilde{d}(k_\parallel)^2, \tag{B2}$$

$$\widetilde{W}_2(k_\parallel) = \int dN_{\rm HI} \, \frac{\partial^2 \mathcal{N}}{\partial\chi \partial N_{\rm HI}} \, \tilde{d}(k_\parallel). \tag{B3}$$

The function $\tilde{d}$ is the Fourier transform of the damping wing absorption profile, $1 - \exp[-\kappa^2 N_{\rm HI}/(\chi - \chi_0)^2]$, χ is the conformal distance, and χ_0 is the value of χ at line-center. We will refer to the first and second terms in the equation for $P_{\rm los}^{\rm DLA}$ as the 1-DLA and 2-DLA terms, respectively. For details, this is analogous to how the dark matter power spectrum in the halo model is calculated, as reviewed in Cooray & Sheth (2002). In addition, $b_{\rm DLA}$ is the bias of such systems (taken to be 2 here), $\partial^2\mathcal{N}/\partial\chi\partial N_{\rm HI}$ is the number of systems per χ per $N_{\rm HI}$, and $\kappa^2 = 5\times10^{-19}$ cm^{-2} Mpc^2 at $z = 3$. At small $k_\parallel$, $\widetilde{W}_2$ becomes equal to the fraction of the continuum that is absorbed by damping wings. Finally, the total effect of damping wings and un-damped Lyα forest absorption on the full 3D flux power spectrum is

$$P_{\rm F}^{\rm all} = \left[b + b_{\rm DLA}\widetilde{W}_2(k_\parallel)\right]^2 \, P_\delta^{\rm lin} + \text{higher order}. \tag{B4}$$

Figure B1 plots the fraction of power that originates from damped systems at $z = 2.5$. These calculations assume that $f(N_{\rm HI}, X) \equiv \partial^2\mathcal{N}/\partial z \partial N_{\rm HI}\sqrt{\Omega_M/(1+z)}$ integrates to 0.1 between 10^{19} and $10^{20.3}$ cm^{-2}, is a power-law with index of -1.2 at $N_{\rm HI} < 10^{20.3}$ cm^{-2}, and with index of -1.8 at higher columns, as motivated in O'Meara et al. (2007) and Prochaska et al. (2010). The function $f(N_{\rm HI}, X)$ is uncertain at the factor of 2-level.

The top panel in Figure B1 shows the ratio of the line-of-sight power in damping wings to $P_{\rm los}$, and the bottom is for the same but for the 3D power. The black thick solid, red thick short-dashed, and blue thick long-dashed curves are respectively the 1-DLA, 2-DLA terms, and cross-power terms. The 1-DLA component is only shown in the top panel because it is a subdominant contribution to the 3D power at the plotted scales owing to the high 3D number density of these systems. The 1-DLA term is a more important contribution than the 2-DLA to $P_{\rm los}^{\rm all}$, and the cross-power term with the forest, $2\, b\, b_{\rm DLA}\, \widetilde{W}_2 P_\delta^{\rm lin}$, is important at higher $k_\parallel$. However, the cross term is the most important contribution to the 3D forest power (bottom panel). The thin dashed curves are the same as the thick curves but only include systems with $N_{\rm HI} > 10^{20}$ cm^{-2} (the contribution that is easiest to remove in pre-processing). Systems with $N_{\rm HI} > 10^{20}$ cm^{-2} contribute most of the 1-DLA power in $P_{\rm los}^{\rm DLA}$ at the smallest k, but contribute little of the 1-DLA power at $k > 0.1$ Mpc^{-1}, of the 2-DLA power in $P_{\rm los}^{\rm all}$ and in $P_{\rm F}^{\rm all}$, or of the cross power.

The morphology and amplitude of the line-of-sight curves in the top panel of Figure B1 are similar to what was found in the numerical calculations of McDonald et al. (2005), who considered $k_\parallel > 0.1$ Mpc^{-1} (see their Fig. 2). [McDonald et al. 2005 found a plateau at the highest k shown in Fig. B1 that likely owes to the higher order terms, such as $P_{\rm los}^{\rm DLA} \star P_{\rm los}$.] Our model ignores the coincidence between the normal Lyα absorption and the damping wing absorption: Damped regions occur where there is already more absorption by un-damped Lyα absorption. McDonald et al. (2005) found that this correlation suppresses the impact of damping wing absorption by a factor that can be as large as 2.

Our toy model for the impact of damping wings provides a couple insights. First, it shows that there is little point of removing the contribution from the highest column systems in 3D analyses of the forest because they contribute a small component of the total power from damping wings. Second, because the 1-DLA term is unimportant in 3D, this reduces the amplitude of the contribution from damped systems relative to 1D analyses, especially at the smallest wavevectors. However, we predict that the contribution from damping wings to $P_{\rm F}$ is still non-negligible. Lastly, in a 3D survey one has the ability to fit for the the damping wing contribution: It should scale at relevant scales as $\mathcal{C}(k_\parallel)\, P_\delta^{\rm lin}(\boldsymbol{k})$, where $\mathcal{C}$ is some function that only depends on $k_\parallel$.